8477231

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

April 10, 2024

THIS IS TO CERTIFY THAT ANNEXED IS A TRUE COPY FROM THE RECORDS OF THIS OFFICE OF THE FILE WRAPPER AND CONTENTS OF:

APPLICATION NUMBER: ***15/836,824***
FILING DATE: ***December 09, 2017***
PATENT NUMBER: ***10686871***
ISSUE DATE: ***June 16, 2020***

Certified by

Performing the Functions and Duties of the
Under Secretary of Commerce
for Intellectual Property
and Director of the United States
Patent and Trademark Office

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 11/30/2014. OMB 0651-0051
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	
Filing Date	
First Named Inventor	Pouya Taaghol
Title	Distributed Software Defined Networking
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001

SIGNATURE of Applicant or Patent Practitioner

Signature	/Christopher J. Horgan/	Date (Optional)	June 2, 2014
Name	Christopher J. Horgan	Registration Number	40,394
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 11/30/2014. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

☑ I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above: 112449

OR

☐ I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

☑ The address associated with the above-mentioned Customer Number

OR

☐ The address associated with Customer Number:

OR

Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

☐ Inventor or Joint Inventor (title not required below)

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)

☑ Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE of Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature		Date (Optional)	
Name	Pouya Tabatabai		
Title	Inventor		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.191, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: ECOMM.AUTH/ECOMM.WTDW
Doc Description: Internet Communications Authorized/Internet Communications Authorization Withdrawn

PTO/SB/439 (11-15)

AUTHORIZATION FOR INTERNET COMMUNICATIONS IN A PATENT APPLICATION OR REQUEST TO WITHDRAW AUTHORIZATION FOR INTERNET COMMUNICATIONS		
	Application No.	
	Filing Date	
	First Named Inventor	TAAGHOL
	Art Unit	
	Examiner Name	
	Practitioner Docket No.	TAAG001CONT1

To: Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

I. To authorize permission for Internet Communications.

☑ Recognizing that Internet communications are not secure, I hereby authorize the USPTO to communicate with the undersigned and practitioners in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application via video conferencing, instant messaging, or electronic mail. I understand that a copy of these communications will be made of record in the application file. (MPEP 502.03)

II. To withdraw authorization for Internet Communications.

☐ The authorization given on _____, to the USPTO to communicate with the undersigned and any practitioner in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application via Internet communications such as video conferencing, instant messaging, or electronic mail is hereby withdrawn. I understand that the withdrawal is effective when approved rather than when received.

I am the

☐ applicant.

☑ attorney or agent of record. Registration number 40,394.

☐ attorney or agent acting under 37 CFR 1.34. Registration number _____.

/Christopher J. Horgan/ — Signature

December 8, 2017 — Date

Christopher J. Horgan — Typed or printed name

408-687-8306 — Telephone Number

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. Juristic entities must be represented by a patent practitioner (see 37 CFR 1.31, which is applicable to any paper filed on or after September 16, 2012, that is presented on behalf of a juristic entity, regardless of application filing date). Submit multiple forms if more than one signature is required, see below*.

☑ * Total of 1 forms are submitted.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

PTO/SB/35 (07-09)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U. S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

NONPUBLICATION REQUEST UNDER 35 U.S.C. 122(b)(2)(B)(i)

First Named Inventor	TAAGHOL
Title	DISTRIBUTED SOFTWARE DEFINED NETWOF
Attorney Docket Number	TAAG001CONT1

I hereby certify that the invention disclosed in the attached application **has not and will not be** the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

I hereby request that the attached application not be published under 35 U.S.C. 122(b).

/Christopher J. Horgan/ — Signature

December 8, 2017 — Date

Christopher J. Horgan — Typed or printed name

40,394 — Registration Number, if applicable

408-687-8306 — Telephone Number

This request must be signed in compliance with 37 CFR 1.33(b) and submitted with the application **upon filing.**

Applicant may rescind this nonpublication request at any time. If applicant rescinds a request that an application not be published under 35 U.S.C. 122(b), the application will be scheduled for publication at eighteen months from the earliest claimed filing date for which a benefit is claimed.

If applicant subsequently files an application directed to the invention disclosed in the attached application in another country, or under a multilateral international agreement, that requires publication of applications eighteen months after filing, the applicant **must** notify the United States Patent and Trademark Office of such filing within forty-five (45) days after the date of the filing of such foreign or international application. **Failure to do so will result in abandonment of this application (35 U.S.C. 122(b)(2)(B)(iii)).**

This collection of information is required by 37 CFR 1.213(a). The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 6 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

PTO/SB/08a (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 1 of 3

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	Taaghol
Art Unit	
Examiner Name	
Attorney Docket Number	TAAGCONT1

U. S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code[2] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
		US-8,528,108	09-2013	Ammer, Gerhard	
		US-8,756,595	06-2014	Arcilla, Andre	
		US-9,226,333	12-2015	Steiner, Moritz	
		US-20090300057	12-2009	Friedman, Nathaniel	
		US-20140241158	08-2014	Anthony, Jr., Bruce	
		US-20110075611	03-2011	Choi, Hyung-Nam	
		US-7,630,706	12-2009	Weigand, David	
		US-20130007254	01-2013	Fries, Robert	
		US-20090113327	04-2009	Pogrebinsky, Vladimir	
		US-20080209491	08-2008	Hasek, Charles	
		US-20140317293	10-2014	Shatzkamer, Kevin	
		US-20140254373	09-2014	Varma, Subir	
		US-20140259012	09-2014	Nandiali, Vishwamitra	
		US-20150173121	06-2015	Miklos, Gyorgy	
		US-20070078988	04-2007	Miloushev, Vladimir	
		US-2009027677	11-2009	Nickolov, Peter	
		US-20020016954	02-2002	Charities, Dietrich	
		US-7114021	09-2006	Seshadri	
		US-8132166	03-2012	DeHaan	

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document Country Code[3]-Number[4]-Kind Code[5] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

PTO/SB/08a (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 2 of 3

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	Taaghol
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001CONT1

U. S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code[2] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
		US-20040107417	06-2004	Chis	
		US-20050246703	11-2005	Ahonen	
		US-20090007091	01-2009	Appiah	
		US-20040237081	11-2004	Homilier, Daniel	
		US-20120300615	11-2012	Kempf, James	
		US-20130160147	06-2013	Draluk, Vadim	
		US-20130263104	10-2013	Baset, Salman	
		US-20140062976	02-2014	Marino, Joseph	
		US-8,650,299	02-2014	Huang, Alex	
		US-8,843,933	09-2014	Holler, Anne Marie	
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document Country Code[3]-Number[4]-Kind Code[5] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

PTO/SB/08b (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 3 of 3

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	TAAGHOL
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001CONT1

NON PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
		Bansal, et al., "OpenRadio: A Programmable Wireless Dataplane", published in HotSDN'12, Aug. 13, 2012, Helsinki, Finland	
		Stallings et al., "Software-Defined Networks and Open Flow" published in The Internet Protocol Journal, Vol. 16, No. 1, March 2012	
		Luo et al., "Software Defined Networking across Distributed Datacenters over Cloud" published 2013 IEEE International Conference on Cloud Computing Technology and Science	
		McKeown et al., "OpenFlow:Enabling Innovation in Campus Networks", March 14, 2008.	
		Li et al., "Toward Software-Defined Cellular Networks", March 14, 2008.	
		"Network Functions Visualization", White Paper published Oct. 22-24, 2012, at the "SDN and OpenFlow World Congress", Damstadt-Germany	
		Yap et al., "OpenRoads: Empowering Research in Mobile Networks" published in 2010 ACM SIGCOMM Computer Communication Review	
		"A Short Introduction To Cloud Platforms " by David Chappell, Aug 2008	

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

2013 IEEE International Conference on Cloud Computing Technology and Science

Software Defined Networking across Distributed Datacenters over Cloud

Mon-Yen Luo and Jun-Yi Chen
Department of Computer Science & Information Engineering,
National Kaohsiung University of Applied Sciences, Taiwan
* myluo@kuas.edu.tw

Abstract- **Cloud computing, in particular its system spread across several data centers, poses a unique set of networking challenges and management problems. This paper presents our design, implementation, and experience in tackling these challenges by using the concepts and technologies of software-defined network (SDN). We define a network abstraction to incorporate the physical and virtual data plane within a datacenter. We propose a new network primitive to address the interconnection issues among multiple datacenters. A policy plane consolidating these mechanisms provides a programmable interface to enable a wide range of policies and new networking mechanisms. We present the performance result on a real deployment over the production networks. The experience and initial performance results demonstrate that the SDN is a promising solution to cloud networking.**

Keywords: Software-Defined Network; Datacenter Network; Cloud Federation; OpenFlow.

I. INTRODUCTION

Cloud computing has emerged as a promising paradigm for cost efficient and pervasive service delivery across data communication networks. While the computing and distributed system issues of cloud technologies have been extensively investigated, fewer attention has been devoted to the unique set of networking challenges to a cloud system. In order to support multi-tenancy and provide more flexible management of resources, the datacenters usually use virtualization technologies (i.e., hypervisors) to abstract the physical resources. The dynamic location of service deployments and virtualization of hardware resources are stressing the network and protocols. Cloud networks must support automated management (creation, deletion, and migration) of enormous numbers of virtual machines (VMs), large numbers of attached physical and virtual devices, and isolated independent subnetworks for multi-tenancy (application components belonging to separate different tenants are collocated on the same physical substrate) [1]. Recent research in cloud networks has mainly focused on networking mechanisms inside the data center [2]. Little attention has been paid to the networking mechanisms for cloud whereby data centers are interconnected through the Internet [2, 3].

Under the support of a project awarded by the National Science Council of Taiwan government, we are currently designing and maintaining a network testbed that interconnects several heterogeneous datacenters of our research partners. Our current implementation and deployment has supported several research projects and applications, including cloud federation [4], international testbed interconnection [5], and network/security research experiments over production networks [6]. We faced some challenges that also exist in some international experimental infrastructures (e.g., GENI [7, 8], and FIRE [9]), cloud federation, or any cloud system spread across multiple datacenters. This paper describes our design, implementation, and experience in tackling these challenges by using the concepts and technologies of software-defined network (SDN). SDN is emerging as a promising solution to make computer networks more programmable, enable innovation in network design, and lower the barrier to deploying new services [10, 11]. The core concepts of SDNs are (i) separating the control (policy) plane from the data (forwarding) plane, (ii) providing well-defined forwarding abstractions (e.g., flow tables in switches) with standard programmatic interfaces (e.g., OpenFlow [12]), (iii) and exposing high-level abstractions to implement a wide range of network policies and new network technologies. Recently, researchers have made quantified arguments that SDN is a viable approach to single datacenter networking [13], enterprise networks [14], or even wide-area traffic management [15].

The major contribution of this paper is toward the design of a system to transform networks among a collection of independent autonomous datacenters into a manageable, flexible, multitenant transport fabric. We begin with an overview of background material on datacenter networking and SDNs, also describing our deployment environment and its networking challenges (Section II). Generally, the challenges arises from the multitenant nature of cloud computing. Each tenant service should be allocated a portion of network resources (we refer to this as a virtual network in this paper) decoupled from the underlying physical equipment. In our deployment, VMs belonging to a virtual network may be hosted either on the hosts within a datacenter or on a collection of physical hosts across different datacenters. We require a system that creates and manages virtual networks on-demand across multiple datacenters. We discuss the design issues and present a high-level overview of our design in Section III, followed by a detailed description in Section IV. We provide the results of performance measurement from the implementation, based on real production networks (Section V). The initial performance data show that our system can create virtual networks across datacenters efficiently. We discuss and conclude our work in Section VI.

II. BACKGROUND AND RELATED WORK

This section first describes background information and related work on data center networks. We then discuss the

978-0-7695-5095-4/13 $31.00 © 2013 IEEE
DOI 10.1109/CloudCom.2013.87

IEEE computer society

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0010

requirements resulting from a system across distributed datacenters and current related solutions.

A. *Datacenter Network*

Data centers exist in various types. Some data centers, for example, use commodity computers running Linux without a hypervisor to process data on a large scale (e.g., Google's datacenters). In this paper, we focus on another popular type of datacenters (i.e., the data centers comprise many racks of commodity servers using hypervisors to support multiple VMs per host). Machines in such datacenters are typically organized into racks and rows under a logical hierarchical network tree [2]. Each rack contains several machines interconnected by a top-of-rack (ToR) switch, which serves as the leaves of the tree and delivers high bandwidth to directly connected hosts. Switches at the leaves have a limited number of high-speed (e.g., 10 GigE) uplinks to one or more network devices that aggregate and transfer packets among the leaf switches. At the root levels of the tree, there are core switches with significant high throughput and switching capacity to relay traffic for inter-row communication. Many proposals and commercial products target on searching for more efficient and low-cost data center network fabrics, such as FatTree [16, 2], and Clos topologies [17, 18, 2]. These topologies face a similar set of challenges, such as improving cost-effectiveness (reducing oversubscription rate), supporting high bisection bandwidth, providing guaranteed bandwidth for each tenant, and increasing fault tolerance.

Much of the research for cloud networking has focused on developing network infrastructure and embedding algorithms optimized for particular topologies. PortLand [19] used a hierarchical pseudo MAC address (PMAC) for all end hosts (or VM) to encode their positions based on a fat tree topology. The edge (ToR) switches, to which a server hosting a VM is connected, perform the rewriting between actual MACs (AMAC) and PMACs. The PMAC design enables efficient and loop-free forwarding (PortLand uses the longest prefix-match), and reduces forwarding states in switches. A fabric manager (a process running on a dedicated machine) maintains a full connectivity matrix of switches and responds to ARP resolution, preventing broadcast flooding. VL2 [20] uses a Clos network topology and implements Valiant Load Balancing (VLB) by sending traffic to random intermediate switches. In VL2, all physical machines and services have application-specific IP addresses (AA), and all switches are assigned a location-specific IP address (LA) with a forwarding mechanism based on OSPF. Similarly, a centralized address manager (AM) maintains the AA-to-LA address mappings. An IP-in-IP encapsulation mechanism is implemented to route on LAs but deliver packets to AAs. NetLord [21] leverages a multipath forwarding scheme termed SPAIN [22] for datacenters using commodity switches. NetLord designs an agent in the hypervisors to encapsulate, route, de-capsulate, and deliver all packets from and to the local VMs. NetLord agents collaborate with central configuration repositories to gather and maintain all the information needed for the encapsulation.

The aforementioned discussion is definitely not a comprehensive list, but represents some important and common characteristics for data center networks. These characteristics led to our design adopting SDN architecture, which will be discussed in Section III.

B. *Issues Related to Distributed Datacenter*

Further challenges exist in a cloud system spread across multiple datacenters, while a virtual network of one particular tenant service may be deployed across (or migrated temporarily to) several datacenters. First, it is very possible that each datacenter has its own addressing scheme and forwarding fabric, which complicates the deployment of a virtual network spread across multiple datacenters. Second, the network infrastructure between each datacenter usually is shared with public Internet services. It is challenging to decouple resources from the public shared Internet facilities to support various tenant services. The traffic among the network interconnection of our testbed can be categorized into three classes: (i) regular Internet traffic, (ii) interconnection traffic between federated clouds (e.g., VM migration, or interaction of service components deployed across multiple datacenters), and (iii) experimental traffic from various research projects. As a result, the third challenge is that we must isolate production and the private cloud network from each other.

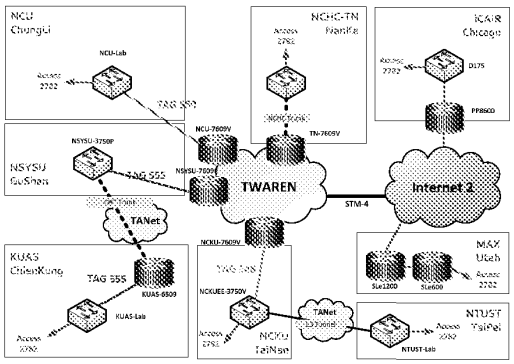


Figure 1: Network Topology of Current Deployment

Let us use current deployment of our testbed to illustrate these issues. As depicted in Figure 1, datacenters in our testbed are majorly interconnected with shared, public layer-3 networks (i.e., the Taiwan Academic Network (TANet) and Taiwan Advanced Research and Education Network (TWAREN)). With supports from iCAIR at Northwestern University and the National Center for High-Performance Computing (NCHC), our testbeds have also been extended to connect with iGENI project [5] in the United States. To isolate the traffic, some of our partner universities (e.g., NCKU, KUAS, NCU) applied some VPLS [23] paths to interlink the private cloud systems. However, VPLS only offers a partial solution to the aforementioned problems. The first reason for this shortcoming is that the use of 12-bit VLAN IDs limits the solution's scalability. At first, in our

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0011

environment, we are only allowed to use the VLAN 2782 to interconnect all sites, which causes another serious scalability and performance problem. The system suffers from issues caused by flooding mechanisms (e.g., ARP broadcast), the huge MAC table size, and the use of Spanning Tree Protocol in a huge layer-2 network [24]. Third, the VLAN, or VPLS, does not have an easy management way to configure the path across several network domains.

We require a system to create and manage simultaneous virtual networks on-demand with arbitrary topologies on loosely coupled distributed datacenters. Some researchers have tried to use IP tunneling techniques and overlay network concepts to extend virtual networks across operational networks, such as VINI [25], VIOLIN [26], and Trellis [27]. In contrast to these systems, our work makes network virtualization coherent in the hypervisor and with the support of NetFPGA, resulting in a good performance and isolation. CloudNet [28] relies on provider-based virtual private networks (VPNs) to extend private network across data centers. The CloudNet is mainly designed for VM live migration, thus, it also implements optimizations for VMs to migrate across the wide area network (WAN). However, the VPN-based mechanisms are mainly point-to-point dedicated connectivity and require a series of interactions and manual operations between network operators and service providers. By contrast, we propose a novel mechanism to extend virtual networks by dynamically organizing and utilizing the VPLS paths (i.e., VLAN tag). The performance result will show that this approach can set up paths to extend a virtual network in only a few seconds while incurring very little overhead in operation.

III. . SYSTEM DESIGN

In this section, we first review the key concepts of SDNs, and then discuss how these principles lead to our design. We then give an overview of our system architecture.

A. *Sofware Defined Networks*

SDN is part of a long history of efforts to make data networks more programmable, which aims to make the network more manageable, enables innovation in network design, and lowers the barrier to deploying new services. In general, SDN has three principal concepts. First, SDN decouples the control plane (which decides how to handle the traffic) and the data plane (which forwards traffic according to decisions that the control plane makes) from each other. Second, SDN logically consolidates the control planes to offer a vision of unified control and management over the states in the network's data-plane elements. Third, SDN provides a well-defined application programming interface (API) to enable the creation of control policies or new applications. OpenFlow [13] is a widely popular instantiation of SDN principles. The OpenFlow defines a data-plane model and a control-plane API by building on technology that commodity switches already supported. The creation of the OpenFlow API was followed quickly by various implementations of controller platforms (e.g., POX [29], Beacon [30], Floodlight [31]) that control the OpenFlow switches and enable new applications. We use OpenFlow in our current implementation and as a concrete instance of SDN throughout the paper. However, we believe that different SDN APIs could be used with our proposed system.

The data-path of an OpenFlow switch is controlled by a table of packet-handling rules. Each rule may contain a pattern (defined by some part of the packet header) that matches a subset of traffic, a list of actions (e.g., drop, flood, forward out a particular interface, modify a header field, send the packet to the controller) on the traffic that matches this rule, a set of counters (to track the number of bytes and packets), and a priority. Upon receiving a packet, an OpenFlow switch identifies the highest-priority matching rule, performs the associated actions, and manipulates the counters. An OpenFlow switch can be programmed (i.e., install forwarding rules) by only one controller (called *master* controller) in either proactive or reactive mode. The controller may proactively install some rules to some of the managed switches. In the reactive mode, an OpenFlow switch sends the first packet from traffic that does not match any rule to a controller, triggering the controller to insert the appropriate flow entries into the table of the OpenFlow switch.

While an important design goal of our system is practically deployable, the SDN architecture provides a practical approach to building a system from the existing commodity switches and real production networks. Based on the observations in Section II, we think that we can exploit SDN architecture to meet our system requirements in the following aspects:

- Many researchers have proposed some important and practical schemes within a single datacenter network. The SDN architecture enables us to extend the network abstraction and design a system architecture that can be compatible with these existing approaches.
- In a multitenant environment like cloud, customers may want the network for their service to be isolated (e.g., a private L2 broadcast domain) from others. This requirement is particularly challenging if such a virtual network is deployed across datacenters. We design a novel mechanism to address this challenge, which can be programmed by an SDN-based architecture.
- The cloud network is still in its infancy, it is therefore important for a management system to provide network control that is flexible enough to support a variety of requirements, and it is extensible to meet a wide range of future needs. The SDN architecture is a promising approach to provide a logically centralized programming environment based on a set of network abstractions.

B. *System Overview*

Our system serves multiple datacenters, each with a number of server racks and network devices. The system architecture can be logically viewed in three layers of software stack, depicted in Figure 2, providing the hooks to our previous work

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0012

or other middleware and serving as a control plane to orchestrate all operations of virtual networks. First, we abstract the data plane of both the physical and virtual switches and synergize them coherently to enable the dynamic configuration and management of virtual networks within a datacenter. We design a software module that resides in the kernel or hypervisor of each host in a datacenter. We also extend a POX [29] controller with some added function to orchestrate the software modules to create a distributed data plane layer.

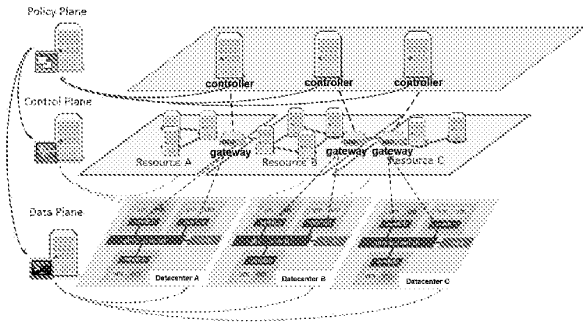


Figure 2: System Overview

The second layer is designed to deal with the interconnection issues among multiple datacenters. The major building blocks of this layer include a gateway switch that is designed based on the OpenFlow protocol. Each datacenter should at least have one gateway system located on the edge of its network to relay traffic across sites. The operation of the gateway system is based on a novel mechanism termed virtualtransit [32]. We assume that each pair of datacenters may set up one or more forwarding paths such as VPLS or MPLS. The basic idea of the virtualtransit is to abstract these paths and dynamically select a suitable path to relay the traffic to the selected datacenter. A distributed control plane layer provides inter-datacenter mechanisms to create and manage virtual private networks spread across public Internet.

The third layer is a logically central control framework containing the global network topology and the policy that determines the behavior of the whole system. The control framework performs route computation on the virtualtransit and explicitly installs flow information to the switches along the chosen route. The control framework consists of multiple controllers that can be replicated for redundancy, and its table can be partitioned for performance scalability. Consolidating all flow decisions to the controller framework, a central service management system can be implemented to enforce a policy throughout the whole system. The system can also be extended further to add new functions and features by simply updating the controller. For instances, our control framework provides interface to allow integration with other control frameworks such as ORCA [33] or our previous control framework [4] for user authentication, or resource registration and discovery.

IV. Implementation and Deployment

Putting all the three layers together, we describe how our system works in more detailed in the following sections. We also describe in some detail our implementation and deployment.

A. *Data Plane Layer*

From the perspective of SDNs, the basic idea of our work in the data plane is to design an appropriate network abstraction to incorporate various addressing and forwarding schemes, so that we can enable a virtual network to be deployed across datacenters in a coherent way. We abstracted the virtual network as a switching element at the core, with a set of network primitives such as port, linking, and policy. In our previous work [3], we implemented a software module that can reside in the hypervisors of host systems to enable the creation and management of virtual networks. In this paper, we further extend our previous work to enable a software-defined networking architecture. We extended the generic control interface of the software module as follows:

- Switching element: This module implemented APIs for creating, removing, and monitoring a virtual switch in a server node. In an SDN architecture, we need a bootstrapping function to make the virtual switch know the controller's location and register itself to the controller. Our system reserves a control channel (VLAN 0 in current implementation) to let switches find and bootstrap connectivity to the controller. Once a connectivity is established, the controller recognizes this switch and aggregates this information to construct the network topology in an information database.
- Virtual port: This module provides functions for adding a virtual port to a VM, removing a virtual port, or disassociating a port from a VM. We also implemented some address assignment mechanisms in this module, which can be customized to support different addressing schemes. We further implement some functions to query and send the link-state information to controller to complete the network topology database.
- Linking: This module provides functions for creating, modifying, and removing connection paths between virtual ports. In our SDN architecture, the controller can configure the virtual switches to perform packet forwarding in three different ways. The first one is called autonomous mode. In current implementation for autonomous mode, we implement two prominent packet-forwarding schemes: SPAIN [22] and VL2 [20]. The autonomous mode is suitable for tenant services that do not require specific topology. The second mode is called proactive mode, which is suitable for tenant service that requires specific topology. For example, recent research has pointed out that the middlebox (e.g., firewalls, network address translators, server load balancers) may further complicate the challenges of managing the datacenter network infrastructure. For some enterprise

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0013

applications, it is necessary to specify the data flow and topology for their networks. We implement functions that are used to create a topology for a given virtual network, including layer-2 mechanisms (e.g., VLAN, VPLS) and layer-3 mechanisms (e.g., IP tunneling, GRE [34]). The third mode is reactive mode, the behavior of which is defined by the controller.

- Policy: This module provides functions for setting a path constraint, such as QoS constraints. These function libraries were designed to allow for further extension to support network management policies such as bandwidth allocation/management mechanisms.

Our mechanisms for intra datacenters do not dictate any particular networking topology or specialized switches. As a proof of concept of system flexibility, we have enabled three datacenter networking architectures with our system. If the intermediate switches are all OpenFlow-enabled switches, our system can support implementations of VL2 or PortLand as in the work [13]. For a datacenter that only uses commodity switches for intermediate switches, our system can use SPAIN by orchestrating the virtual switches and physical switches.

B. Control-Plane: Inter-datacenter Mechanisms

In our latest work [32], we have proposed a novel mechanism to address the interconnection issues among multiple datacenters. Our system does not require any modification to the switches in the datacenters. We add only some OpenFlow switches (OFS for short) that serve as a gateway system for the interconnection of datacenters. The data plane of OFS remains, and only processes packets according to the flow table, which is configured by the controller. The concept of virtualtransit is novel and promising because we can abstract the interconnection problem by adding a new network primitive that a single system can control dynamically. In this work, we further extend and generalize this mechanism by introducing the SDN architecture. Based on the network topology information, the controller framework manipulates a flow table with two parts (i.e., ARP table and transit table) to execute the virtualtransit. We use Figure 3 to show how the mechanism could operate in a real deployment. When a virtual network must be deployed across datacenters, an offline algorithm is invoked to compute the transit paths and configure the related information into the transit table. The major challenge of the inter-cloud mechanism is how to relay necessary information and isolate the local traffic of each datacenter. The controller framework manipulates the ARP table to address this issue. The reader is referred to [32] for further details on the operation of virtualtransit.

As a proof of concept, we use one single controller in the work described in [32]. While our measurements in Section V show that a single controller system may be able to manage thousands of machines, it also suggests that the network latency is a considerable concern while the gateway system needs to connect the controller across wide area network. As a result, multiple controllers forming a logically single control platform is desirable to provide high-availability or to scale to more datacenters over wide area networks. Our design of a distributed controller takes advantage of two properties of system behaviors and workload characteristics. First, the flow tables can easily be partitioned to several locations because the physical network topology (e.g., the possible transit path) is the same for all replicas, and its placement and binding can be handled in a purely distributed fashion without a strong consistency requirement on table content. Second, similar to the observation in an enterprise network [14], changes to the network state happen on very slow time scales relative to network flow arrivals. Thus, a standard replication strategy is enough to ensure the correctness of forwarding on each site and maintain consistency in global. As a result, we now use a full replication strategy for the flow table on each controller. Only some events (such as topology relocation, link failure, and change of transit path) result in a change of the controller state and incur overhead for consistency management. We expect future exploration on state partition or fault tolerance to make the system more efficient and robust

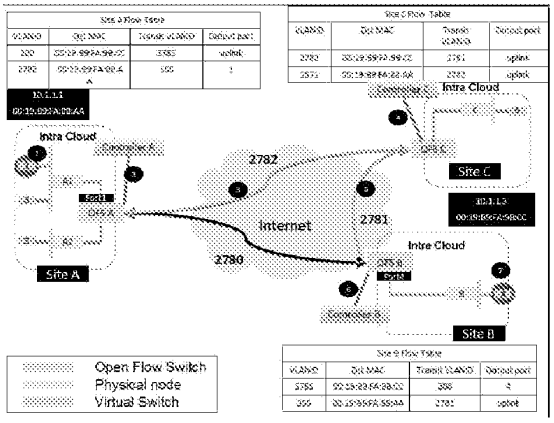


Figure 3: Operation of Inter cloud Mechanisms

C. Policy Plane

Based on the aforementioned two layers, we now can provide a logically centralized control plane that offers the programmability required to dynamically configure and manage the whole network. Due to the space limitation, we only discuss how a policy can be enforced throughout the system as a use case. We leave the comprehensive discussion to our future publications. Let us use Figure 3 to illustrate this use case. Figure 3 represents a real deployment example in our federated datacenter platform. In this example, five VMs (VM1 to VM5) belonging to a single tenant were allocated to different sites. System administrator can use the policy plane to specify the topology of this virtual network and the management policies for this network. The virtual network is then deployed to the network environment as depicted in Figure 3. The subnetwork in

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0014

site A used VLAN 100 for intra cloud communication, and the sub-network in site C used VLAN 200. Such a design is scalable because each datacenter may have complete and independent VLAN space, and an overlapping VLAN tag space is allowable. In contrast, the VPLS or VLAN in WAN environment is more expensive to deploy and manage. The virtualtransit is designed to solve this problem. A path computation mechanism is invoked to select a suitable path. In the example illustrated in Figure 3, that virtual network eventually was deployed using VPLS path 2782 (from Site A to Site C) and VPLS 2781 (from Site C to Site B).

Based on the topology requirement and path selection, the policy plane then parses the topology-related information and configures the ARP table and transit table. When the ARP query for VM4 is forwarded to OFSA, OFSA will consult the ARP table for the mapping, and then serve as the ARP proxy to respond to the ARP query. Such a design prevents ARP flooding and makes the system scalable and efficient. The module also updates the tables when one VM is migrated or destroyed because the VM creation, migration, and destruction are controlled by the management system or cloud manager. This confirms that an SDN architecture can support VM live migration efficiently by controllers, which install flow entries in all switches along the new path and inactivate (or remove) flow entries from the switches in the previous path upon receiving notification that the VM is moving.

When the first packet (after ARP query and reply) of the traffic flow from VM1 to VM4 arrives at the OFSA, the OFSA sends a control packet to the controller because it is missing in the local flow table. The information (OSF ID, VLAN ID, and Destination MAC) from the control packet is used as keys to search the mapping entry in the transit table. If a mapping entry is found, the transit VLAN ID and output port will be sent back to the OFS as an action that instructs the OpenFlow switch to forward a packet. In our example, the packet from VM1 to VM4 will be sent to OFSC, and the VLAN tag will be translated from 100 to 2782. Then the packet is sent to OFSB following a similar process (i.e., translate the VLAN 2782 to 2781), and the packet is sent to Site B by translating 2781 to 200. The gateway system manipulates the two tables to conduct VLAN tag translation between the network domains. With the support of NetFPGA [35], the management of VLAN tags translation is maintained in the control plane and leaves the data-plane processing and tag-translation processing in the line rate.

V. Performance Evaluation

We have deployed the proposed system in operational networks entrusting our design and work in production networks. We implemented the gateway switch based on the NetFPGA programmable switch platform that has four Gigabit Ethernet ports, a Xilinx Virtex-II FPGA, and 4 MB of SRAM. We implemented the controller on a standard Linux PC (2.50GHz Intel Xeon CPU with a 4GB RAM and 1GE network interface). Due to space limitations and the scope of this paper, we focus on the performance results regarding the SDN approach itself in the following discussion. We leave the performance data regarding the achieved throughput of forwarding, the overhead introduced by the proposed mechanisms, and benefits of flexible virtual network deployment to our other publications [3, 32].

A. *Flexibility*

The major goal of the proposed system is to provide network control that is flexible enough to meet a wide range of needs in a heterogeneous environment interconnected by datacenters. Thus, it is important to verify the ability of our system to flexibly specify and instantiate a variety of virtual networks in the cloud comprising multiple datacenters. With the proposed system in this paper, we can easily set up virtual networks as shown in Figure 3 for proof of concept and performance measurement in real production networks. There are three groups of physical machines; the first two (Sites A and C in Figure 3) are located in KUAS, and the third (Site B) is located at another university (NCKU) with eight routing hops between the two universities. The first two groups are allocated in different datacenters and assigned different IP subnets. The traffic between these groups is sent across the WAN environment. All machines use the 2.50GHz Intel Xeon CPU (E5420) with a 4GB RAM and 1GE network interface.

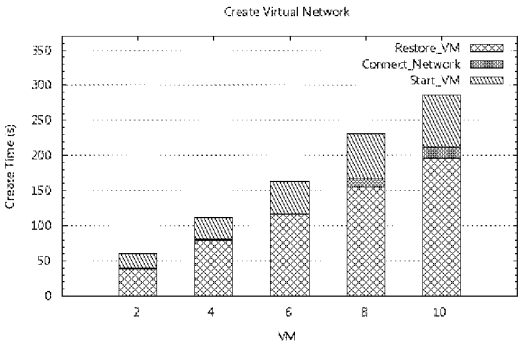


Figure 4: Performance Result of Creation of Virtual Networks

With our system support, we can create virtual networks across three datacenters in a single command. Figure 4 shows the completion time of five virtual networks (2, 4, 6, 8, and 10 nodes in each virtual network, respectively). The completion time is the time from when the network requirement is specified and submitted to when the whole connectivity of this network is established and the traffic can flow successfully (first ping can be responded by each node). The time required in creation of virtual network is only the part labeled with *connect_netwok* in the Figure 4. We can see that our system can deploy virtual networks efficiently. The largest creation time of a virtual network is less than three seconds, in contrast to the previous approaches proposed in [28] in which 27 seconds were needed using VPN.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0015

B. *Scalability*

The size of flow tables in control framework is an important factor for system scalability [13], so we estimate the size of the tables needed by controllers. To emulate a larger environment, we wrote a simulation program based on topology and empirical statistics from our deployment and other publications [15,19,20,21]. In this simulation, we assume a network environment consisting of four datacenters. We perform our analysis by varying the size (number of servers) of each datacenter and the ratio of the cross-site traffic. For example, in Experiment 4 of Figure 5, each datacenter has 50K VMs, and 20% of these VMs have cross-site connections to the nodes in other datacenter. Our simulation is a multithread program, and each thread undergoes the following steps in each experiment: (1) randomly select a VM of one particular datacenter, (2) determine the size of the virtual network based on statistical distribution from system logs on real deployment, (3) randomly select the VMs to make the virtual network and meet its topology requirement, (4) construct the tables required in the controller according the choices made in the previous three steps, and (5) maintain the network for a period time based on the traffic characteristic, then dissolve the network and remove related entries in the tables. The numbers and the required thread are determined by the numbers of hosts and the ratio specified in each experiment.

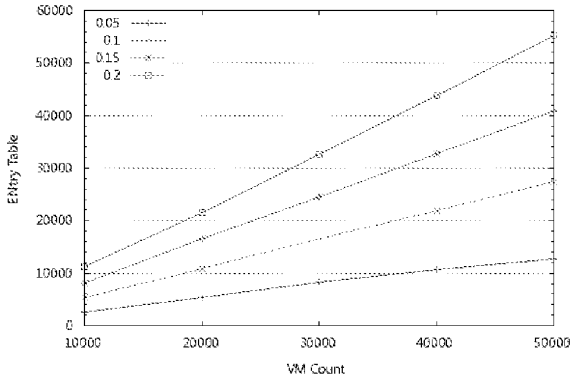


Figure 5: Performance Result of Latency

We realize that other datacenters may have different topologies and the reference data characteristics may not be comprehensive, but our goal here is to provide an initial estimate of our design to cope with large-scale deployment across multiple datacenters. We recorded the maximum number of entries in the controller's table for each experiment in Figure 5. In the largest topology considered in our simulation (Experiment 4), we had a maximum numbers of table entries (55K entries). This initial result shows that the table size is not prohibitively large for a mid-sized cloud with multiple datacenters. A single controller could comfortably handle 10K flow requests per second with such a table size. Please note that not every new flow needs to query the controller; only the new flow whose destination is another datacenter needs to be sent to the controller. However, we observe that the network latency of flow setup time is a significant factor. We perform another measurement in our testbed with some real applications to verify this issue. In the first configuration of this measurement, there is only one controller (located in NCKU, see the Figure 1). As a result, the gateway switch must connect the controller across WAN for flow setup. In the second configuration, we set up another controller in KUAS, so that a portion of flow queries can be handled by the second controller locally. The distributions of response time for flow setup are shown in Figure 6, which is taken in samples over the course of 7 days. The data show that the multiple controllers can improve the flow setup time significantly.

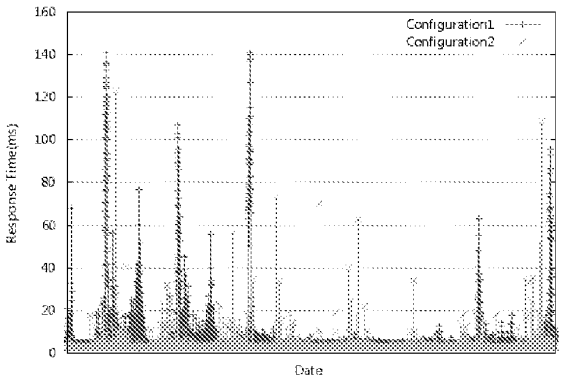


Figure 6: Comparison of Flow Setup Time in Two Configurations

VI. CONCLUSION

We have demonstrated how to leverage the software-defined network concepts and technologies to design and implement an integrated system to address some networking challenges in cloud computing. While the networking mechanism in cloud system is still rapidly evolving, the programmability and extensibility of an SDN-based architecture provide a promising solution to solve pressing problems in cloud networking. Our experience and initial performance results also reveal some important issues for further research. First, although the idea of virtualtransit is promising to solve the interconnection issues among datacenter federation, further investigation are needed on the issues about policy negotiation, security mechanism, and path selection algorithm. Second, more research efforts should be devoted to distributed system issues involved in the design of a control plane scalable to a large network deployment.

ACKNOWLEDGEMENTS

This research was financially supported by the National Science Council of Taiwan, under grants No. NSC 100-2218-E-151-003-MY3 and NSC 101-2622-E-151-026-CC3. The authors also thank NCKU and NCHC for their support on network facilities and WAN configurations.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0016

REFERENCES

[1] C. J. S. Decusatis, A. Carranza, and C. M. Decusatis, "Communication within clouds: open standards and proprietary protocols for data center networking," Communications Magazine, IEEE, vol. 50, pp. 26-33, 2012.

[2] M. F. Bari, R. Boutaba, R. Esteves, L. Z. Granville, M. Podlesny, M. G. Rabbani, Z. Qi, and M. F. Zhani, "Data Center Network Virtualization: A Survey," Communications Surveys & Tutorials, IEEE, vol. 15, pp. 909-928, 2013.

[3] M.-Y. Luo and J.-Y. Chen, "Towards Network Virtualization Management for Federated Cloud Systems", IEEE 6th International Conference on Cloud Computing, June 27-July 2, 2013.

[4] M.-Y. Luo and S.-W. Lin, "From Monolithic Systems to a Federated E-Learning Cloud System," in IEEE International Conference on Cloud Engineering, San Francisco, California, USA, 2013.

[5] J. Mambretti, J. Chen, F. Yeh, T.-L. Liu, M.-Y. Luo, C.-S. Yang, R. v. d. Pol, A. Barczyk, F. Dijkstra, and G. v. Malensteinz, "OpenFlow Services for Science: An International Experimental Research Network Demonstrating Multi-Domain Automatic Network Topology Discovery, Direct Dynamic Path Provisioning Using Edge Signaling and Control, Integration With Multipathing Using MPTCP," in SCinet Research Sandbox at International Conference for high performance Computing, Networking, Storage, and Analysis, 2012.

[6] M.-Y. Luo, J. Chen, J. Mambretti, F. Yeh, and C.-S. Yang, "Network Virtualization Implementation Over Global Research Production Networks," Journal of Internet Technology, 2014.

[7] GENI: Global Environment for Network Innovations, http://www.geni.net

[8] L. Peterson, T. Anderson, D. Blumenthal, D. Casey, D. Clark, D. Estrin, J. Evans, D. Raychaudhuri, M. Reiter, J. Rexford, S. Shenker, and J. Wroclawski, "GENI Design Principles," Computer, vol. 39, pp. 102-105, 2006.

[9] FIRE: Future Internet Research and Experimentation, http://cordis.europa.eu/fp7/ict/fire/

[10] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A clean slate 4D approach to network control and management," ACM SIGCOMM Computer Communication Review, vol. 35, pp. 41-54, 2005.

[11] N. Gude, T. Koponen, J. Pettit, B. Pfaff, M. Casado, N. McKeown, and S. Shenker, "NOX: towards an operating system for networks," ACM SIGCOMM Computer Communication Review, vol. 38, pp. 105-110, 2008.

[12] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: enabling innovation in campus networks," ACM SIGCOMM Computer Communication Review, vol. 38, pp. 69-74, 2008.

[13] A. Tavakoli, M. Casado, T. Koponen, and S. Shenker, "Applying NOX to the Datacenter," in HotNets, 2009.

[14] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker, "Ethane: Taking control of the enterprise," in ACM SIGCOMM Computer Communication Review, 2007, pp. 1-12.

[15] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, and M. Zhu, "B4: Experience with a globally-deployed software defined WAN," in Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM, pp. 3-14, 2013.

[16] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in ACM SIGCOMM Computer Communication Review, 2008, pp. 63-74.

[17] C. Clos, "A study of non-blocking switching networks," Bell System Technical Journal, vol. 32, pp. 406-424, 1953.

[18] C. Kim, M. Caesar, and J. Rexford, "Floodless in seattle: a scalable ethernet architecture for large enterprises," in ACM SIGCOMM Computer Communication Review, pp. 3-14, 2008.

[19] R. Niranjan Mysore, A. Pamboris, N. Farrington, N. Huang, P. Miri, S. Radhakrishnan, V. Subramanya, and A. Vahdat, "PortLand: a scalable fault-tolerant layer 2 data center network fabric," in ACM SIGCOMM Computer Communication Review, pp. 39-50, 2009.

[20] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta, "VL2: a scalable and flexible data center network," in ACM SIGCOMM Computer Communication Review, pp. 51-62, 2009.

[21] J. Mudigonda, P. Yalagandula, J. Mogul, B. Stiekes, and Y. Pouffary, "NetLord: a scalable multi-tenant network architecture for virtualized datacenters," SIGCOMM-Computer Communication Review, vol. 41, p. 62, 2011.

[22] J. Mudigonda, P. Yalagandula, M. Al-Fares, and J. C. Mogul, "SPAIN: COTS Data-Center Ethernet for Multipathing over Arbitrary Topologies," in USENIX Symposium on Networked Systems Design and Implementation (NSDI), pp. 265-280, 2010.

[23] M. Lasserre and V. Kompella, "Virtual private LAN service (VPLS) using label distribution protocol (LDP) signaling," RFC 4762, January 2007.

[24] C. Kim and J. Rexford, "Revisiting Ethernet: Plug-and-play made scalable and efficient," in Local & Metropolitan Area Networks, 2007. LANMAN 2007. 15th IEEE Workshop on, 2007, pp. 163-169.

[25] Andy C. Bavier, Nick Feamster, Mark Huang, Larry L. Peterson and Jennifer Rexford. In VINI Veritas: Realistic and Controlled Network Experimentation. In Proceedings of the SIGCOMM Conference, pages 3-14, Piza, Italy, September 2006.

[26] Xuxian Jiang and Dongyan Xu. Violin: Virtual internetworking on overlay infrastructure. In International Symposium on Parallel and Distributed Processing and Applications (ISPA) 2004, 2004.

[27] S. Bhatia, M. Motiwala, W. Muhlbauer, Y. Mundada, V. Valancius, A. Bavier, N. Feamster, L. Peterson, and J. Rexford. "Trellis: A platform for building flexible, fast virtual networks on commodity hardware". Proc. Workshop on Real Overlays and Distributed Systems, 2008.

[28] T. Wood, K. K. Ramakrishnan, P. Shenoy, and J. Van der Merwe, "CloudNet: dynamic pooling of cloud resources by live WAN migration of virtual machines," in ACM SIGPLAN Notices, pp. 121-132, 2011.

[29] POX controller. http://www.noxrepo.org/pox/

[30] Beacon: http://beaconcontroller.net/.

[31] Floodlight OpenFlow Controller. http://floodlight.openflowhub.org/.

[32] M.-Y. Luo and J.-Y. Chen, "VirtualTransits: a Platform for Network Virtualization across Data Centers," 2nd IEEE International Conference on Cloud Networking, November 2013.

[33] J. Chase, I. Constandache, A. Demberel, L. Grit, V. Marupadi, M. Sayler, and A. Yumerefendi, "Controlling dynamic guests in a virtual computing utility," In International Conference on the Virtual Computing Initiative (ICVCI 2008), May 2008.

[34] D. Farinacci, T. Li, S. Hanks, D. Meyer, and P. Traina. GenericRouting Encapsulation (GRE). Internet Engineering Task Force, March 2000. RFC 2784.

[35] NetFPGA, http://netfpga.org/.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0017

Log In My Cisco
Products & Services Support How to Buy Training & Events Partners

Software-Defined Networks and OpenFlow - The Internet Protocol Journal, Volume 16, No. 1

HOME
ABOUT CISCO
PRESS
THE INTERNET PROTOCOL JOURNAL
ISSUES
VOLUME 16, NUMBER 1, MARCH 2013
From the Editor
SDN and OpenFlow
Address Authentication
WCIT Report
Letters to the Editor
Book Review
Fragments
Call for Papers

by William Stallings

A network organizing technique that has come to recent prominence is the *Software-Defined Network* (SDN) [1]. In essence, an SDN separates the data and control functions of networking devices, such as routers, packet switches, and LAN switches, with a well-defined *Application Programming Interface* (API) between the two. In contrast, in most large enterprise networks, routers and other network devices encompass both data and control functions, making it difficult to adjust the network infrastructure and operation to large-scale addition of end systems, virtual machines, and virtual networks. In this article we examine the characteristics of an SDN, and then describe the *OpenFlow* specification, which is becoming the standard way of implementing an SDN.

Evolving Network Requirements

Before looking in more detail at SDNs, let us examine the evolving network requirements that lead to a demand for a flexible, response approach to controlling traffic flows within a network or the Internet.

One key leading factor is the increasingly widespread use of *Server Virtualization*. In essence, server virtualization masks server resources, including the number and identity of individual physical servers, processors, and operating systems, from server users. This masking makes it possible to partition a single machine into multiple, independent servers, conserving hardware resources. It also makes it possible to migrate a server quickly from one machine to another for load balancing or for dynamic switchover in the case of machine failure. Server virtualization has become a central element in dealing with "big data" applications and in implementing cloud computing infrastructures. But it creates problems with traditional network architectures (for example, refer to [2]). One problem is configuring *Virtual LANs* (VLANs). Network managers need to make sure the VLAN used by the *Virtual Machine* is assigned to the same switch port as the physical server running the virtual machine. But with the virtual machine being movable, it is necessary to reconfigure the VLAN every time that a virtual server is moved. In general terms, to match the flexibility of server virtualization, the network manager needs to be able to dynamically add, drop, and change network resources and profiles. This process is difficult to do with conventional network switches, in which the control logic for each switch is co-located with the switching logic.

Another effect of server virtualization is that traffic flows differ substantially from the traditional client-server model. Typically, there is a considerable amount of traffic among virtual servers, for such purposes as maintaining consistent images of the database and invoking security functions such as access control. These server-to-server flows change in location and intensity over time, demanding a flexible approach to managing network resources.

Another factor leading to the need for rapid response in allocating network resources is the increasing use by employees of mobile devices such as smartphones, tablets, and notebooks to access enterprise resources. Network managers must be able to respond to rapidly changing resource, *Quality of Service* (QoS), and security requirements.

Existing network infrastructures can respond to changing requirements for the management of traffic flows, providing differentiated QoS levels and security levels for individual flows, but the process can be very time-consuming if the enterprise network is large and/or involves network devices from multiple vendors. The network manager must configure each vendor's equipment separately, and adjust performance and security parameters on a per-session, per-application basis. In a large enterprise, every time a new virtual machine is brought up, it can take hours or even days for network managers to do the necessary reconfiguration [3].

This state of affairs has been compared to the mainframe era of computing [4]. In the era of the mainframe, applications, the operating system, and the hardware were vertically integrated and provided by a single vendor. All of these ingredients were proprietary and closed, leading to slow innovation. Today, most computer platforms use the x86 instruction set, and a variety of operating systems (Windows, Linux, or Mac OS) run on top of the hardware. The OS provides APIs that enable outside providers to develop applications, leading to rapid innovation and deployment. In a similar fashion, commercial networking devices have proprietary features and specialized control planes and hardware, all vertically integrated on the switch. As will be seen, the SDN architecture and the OpenFlow standard provide an open architecture in which control functions are separated from the network device and placed in accessible control servers. This setup enables the underlying infrastructure to be abstracted for applications and network services, enabling the network to be treated as a logical entity.

SDN Architecture

Figure 1 illustrates the logical structure of an SDN. A central controller performs all complex functions, including routing, naming, policy declaration, and security checks. This plane constitutes the *SDN Control Plane*, and consists of one or more SDN servers.

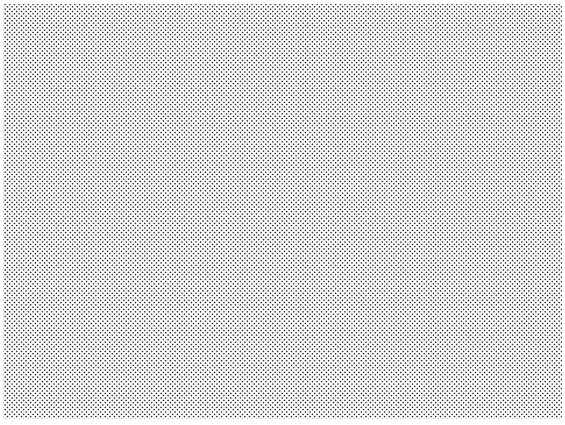


Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0018

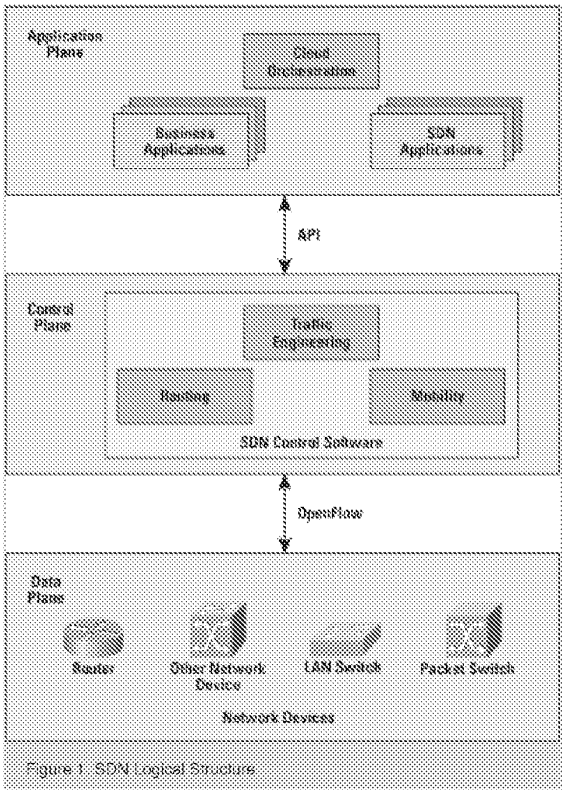


Figure 1: SDN Logical Structure

The *SDN Controller* defines the data flows that occur in the *SDN Data Plane*. Each flow through the network must first get permission from the controller, which verifies that the communication is permissible by the network policy. If the controller allows a flow, it computes a route for the flow to take, and adds an entry for that flow in each of the switches along the path. With all complex functions subsumed by the controller, switches simply manage flow tables whose entries can be populated only by the controller. Communication between the controller and the switches uses a standardized protocol and API. Most commonly this interface is the OpenFlow specification, discussed subsequently.

The SDN architecture is remarkably flexible: it can operate with different types of switches and at different protocol layers. SDN controllers and switches can be implemented for Ethernet switches (Layer 2), Internet routers (Layer 3), transport (Layer 4) switching, or application layer switching and routing. SDN relies on the common functions found on networking devices, which essentially involve forwarding packets based on some form of flow definition.

In an SDN architecture, a switch performs the following functions:

- The switch encapsulates and forwards the first packet of a flow to an SDN controller, enabling the controller to decide whether the flow should be added to the switch flow table.
- The switch forwards incoming packets out the appropriate port based on the flow table. The flow table may include priority information dictated by the controller.
- The switch can drop packets on a particular flow, temporarily or permanently, as dictated by the controller. Packet dropping can be used for security purposes, curbing Denial-of-Service (DoS) attacks or traffic management requirements.

In simple terms, the SDN controller manages the forwarding state of the switches in the SDN. This management is done through a vendor-neutral API that allows the controller to address a wide variety of operator requirements without changing any of the lower-level aspects of the network, including topology.

With the decoupling of the control and data planes, SDN enables applications to deal with a single abstracted network device without concern for the details of how the device operates. Network applications see a single API to the controller. Thus it is possible to quickly create and deploy new applications to orchestrate network traffic flow to meet specific enterprise requirements for performance or security.

SDN Domains

In a large enterprise network, the deployment of a single controller to manage all network devices would prove unwieldy or undesirable. A more likely scenario is that the operator of a large enterprise or carrier network divides the whole network into numerous nonoverlapping SDN domains as shown in Figure 2.

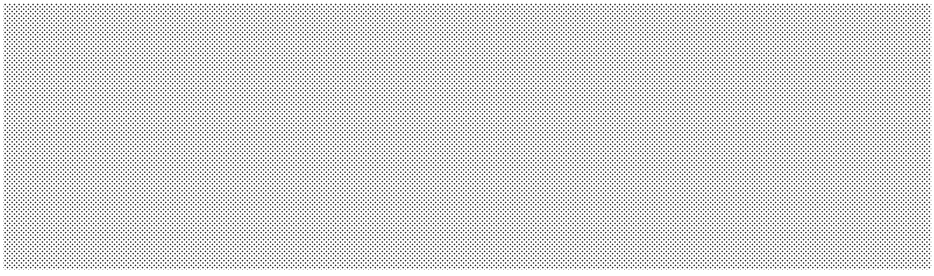


Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0019

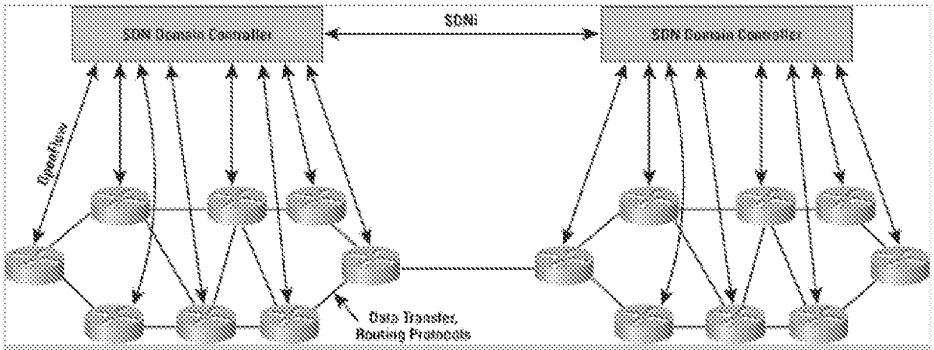


Figure 2: SDN Domain Structure

Reasons for using SDN domains include the following:

- *Scalability:* The number of devices an SDN controller can feasibly manage is limited. Thus, a reasonably large network may need to deploy multiple SDN controllers.
- *Privacy:* A carrier may choose to implement different privacy policies in different SDN domains. For example, an SDN domain may be dedicated to a set of customers who implement their own highly customized privacy policies, requiring that some networking information in this domain (for example, network topology) not be disclosed to an external entity.
- *Incremental deployment:* A carrier's network may consist of portions of traditional and newer infrastructure. Dividing the network into multiple, individually manageable SDN domains allows for flexible incremental deployment.

The existence of multiple domains creates a requirement for individual controllers to communicate with each other via a standardized protocol to exchange routing information. The IETF is currently working on developing a protocol, called *SDNi,* for "interfacing SDN Domain Controllers" [5]. SDNi functions include:

- Coordinate flow setup originated by applications containing information such as path requirement, QoS, and service-level agreements across multiple SDN domains.
- Exchange reachability information to facilitate inter-SDN routing. This information exchange will allow a single flow to traverse multiple SDNs and have each controller select the most appropriate path when multiple such paths are available.

The message types for SDNi tentatively include the following:

- Reachability update
- Flow setup/tear-down/update request (including application capability requirements such as QoS, data rate, latency etc.)
- Capability update (including network-related capabilities such as data rate and QoS, and system and software capabilities available inside the domain)

OpenFlow

To turn the concept of SND into practical implementation, two requirements must be met. First, there must be a common logical architecture in all switches, routers, and other network devices to be managed by an SDN controller. This logical architecture may be implemented in different ways on different vendor equipment and in different types of network devices, so long as the SDN controller sees a uniform logical switch function. Second, a standard, secure protocol is needed between the SDN controller and the network device.

Both of these requirements are addressed by *OpenFlow,* which is both a protocol between SDN controllers and network devices, as well as a specification of the logical structure of the network switch functions [6, 7]. OpenFlow is defined in the *OpenFlow Switch Specification,* published by the *Open Networking Foundation* (ONF). ONF is a consortium of software providers, content delivery networks, and networking equipment vendors whose purpose is to promote software-defined networking.

This discussion is based on the current OpenFlow specification, Version 1.3.0, June 25, 2012 [8]. The original specification, 1.0, was developed at Stanford University and was widely implemented. OpenFlow 1.2 was the first release from ONF after inheriting the project from Stanford. OpenFlow 1.3 significantly expands the functions of the specification. Version 1.3 is likely to become the stable base upon which future commercial implementations for OpenFlow will be built. ONF intends for this version to be a stable target for chip and software vendors, so little if any change is planned for the foreseeable future [9].

Logical Switch Architecture

Figure 3 illustrates the basic structure of the OpenFlow environment. An SDN controller communicates with OpenFlow-compatible switches using the OpenFlow protocol running over the *Secure Sockets Layer* (SSL). Each switch connects to other OpenFlow switches and, possibly, to end-user devices that are the sources and destinations of packet flows. Within each switch, a series of tables—typically implemented in hardware or firmware—are used to manage the flows of packets through the switch.

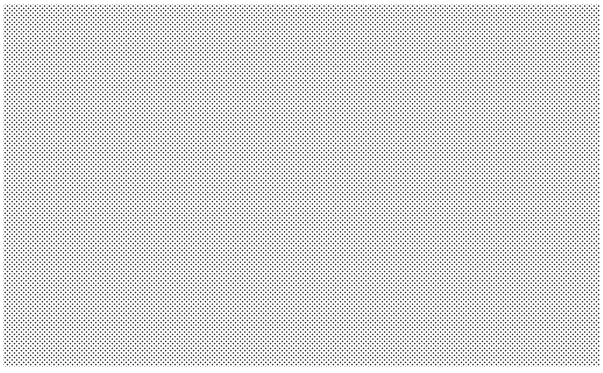


Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0020

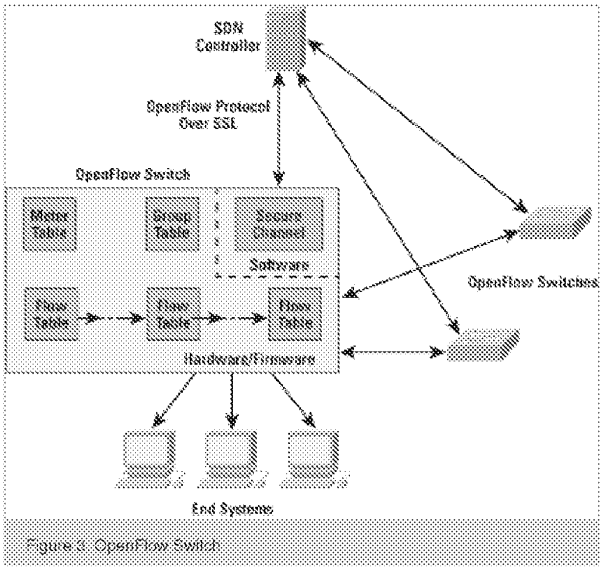


Figure 3: OpenFlow Switch

The OpenFlow specification defines three types of tables in the logical switch architecture. A *Flow Table* matches incoming packets to a particular flow and specifies the functions that are to be performed on the packets. There may be multiple flow tables that operate in a pipeline fashion, as explained subsequently. A flow table may direct a flow to a *Group Table,* which may trigger a variety of actions that affect one or more flows. A *Meter Table* can trigger a variety of performance-related actions on a flow.

Before proceeding, it is helpful to define what the term flow means. Curiously, this term is not defined in the OpenFlow specification, nor is there an attempt to define it in virtually all of the literature on OpenFlow. In general terms, a flow is a sequence of packets traversing a network that share a set of header field values. For example, a flow could consist of all packets with the same source and destination IP addresses, or all packets with the same VLAN identifier. We provide a more specific definition subsequently.

Flow Table Components

The basic building block of the logical switch architecture is the flow table. Each packet that enters a switch passes through one or more flow tables. Each flow table contains entries consisting of six components:

- *Match Fields:* Used to select packets that match the values in the fields.
- *Priority:* Relative priority of table entries.
- *Counters:* Updated for matching packets. The OpenFlow specification defines a variety of timers. Examples include the number of received bytes and packets per port, per flow table, and per flow-table entry; number of dropped packets; and duration of a flow.
- *Instructions:* Actions to be taken if a match occurs.
- *Timeouts:* Maximum amount of idle time before a flow is expired by the switch.
- *Cookie:* Opaque data value chosen by the controller. May be used by the controller to filter flow statistics, flow modification, and flow deletion; not used when processing packets.

A flow table may include a table-miss flow entry, which renders all Match Fields wildcards (every field is a match regardless of value) and has the lowest priority (priority 0). The Match Fields component of a table entry consists of the following required fields:

- *Ingress Port:* The identifier of the port on the switch where the packet arrived. It may be a physical port or a switch-defined virtual port.
- *Ethernet Source and Destination Addresses:* Each entry can be an exact address, a bitmasked value for which only some of the address bits are checked, or a wildcard value (match any value).
- *IPv4 or IPv6 Protocol Number:* A protocol number value, indicating the next header in the packet.
- *IPv4 or IPv6 Source Address and Destination Address:* Each entry can be an exact address, a bitmasked value, a subnet mask value, or a wildcard value.
- *TCP Source and Destination Ports:* Exact match or wildcard value.
- *User Datagram Protocol (UDP) Source and Destination Ports:* Exact match or wildcard value.

The preceding match fields must be supported by any OpenFlow-compliant switch. The following fields may be optionally supported:

- *Physical Port*: Used to designate underlying physical port when packet is received on a logical port.
- *Metadata:* Additional information that can be passed from one table to another during the processing of a packet. Its use is discussed subsequently.
- *Ethernet Type:* Ethernet Type field.
- *VLAN ID and VLAN User Priority:* Fields in the IEEE 802.1Q Virtual LAN header.
- *IPv4 or IPv6 DS and ECN:* Differentiated Services and Explicit Congestion Notification fields.
- *Stream Control Transmission Protocol (SCTP) Source and Destination Ports:* Exact match or wildcard value.
- *Internet Control Message Protocol (ICMP) Type and Code Fields:* Exact match or wildcard value.
- *Address Resolution Protocol (ARP) Opcode:* Exact match in Ether-net Type field.
- *Source and Target IPv4 Addresses in Address Resolution Protocol (ARP) Payload:* Can be an exact address, a bitmasked value, a subnet mask value, or a wildcard value.
- *IPv6 Flow Label:* Exact match or wildcard.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0021

- *ICMPv6 Type and Code fields:* Exact match or wildcard value.
- *IPv6 Neighbor Discovery Target Address:* In an IPv6 Neighbor Discovery message.
- *IPv6 Neighbor Discovery Source and Target Addresses:* Link-layer address options in an IPv6 Neighbor Discovery message.
- *Multiprotocol Label Switching (MPLS) Label Value, Traffic Class, and Bottom of Stack (BoS):* Fields in the top label of an MPLS label stack.

Thus, OpenFlow can be used with network traffic involving a variety of protocols and network services. Note that at the MAC/link layer, only Ethernet is supported. Thus, OpenFlow as currently defined cannot control Layer 2 traffic over wireless networks.

We can now offer a more precise definition of the term *flow*. From the point of view of an individual switch, a flow is a sequence of packets that matches a specific entry in a flow table. The definition is packet-oriented, in the sense that it is a function of the values of header fields of the packets that constitute the flow, and not a function of the path they follow through the network. A combination of flow entries on multiple switches defines a flow that is bound to a specific path.

The *instructions component* of a table entry consists of a set of instructions that are executed if the packet matches the entry. Before describing the types of instructions, we need to define the terms "Action" and "Action Set." Actions describe packet forwarding, packet modification, and group table processing operations. The OpenFlow specification includes the following actions:

- *Output:* Forward packet to specified port.
- *Set-Queue:* Sets the queue ID for a packet. When the packet is forwarded to a port using the output action, the queue id determines which queue attached to this port is used for scheduling and forwarding the packet. Forwarding behavior is dictated by the configuration of the queue and is used to provide basic QoS support.
- *Group:* Process packet through specified group.
- *Push-Tag/Pop-Tag:* Push or pop a tag field for a VLAN or MPLS packet.
- *Set-Field:* The various Set-Field actions are identified by their field type; they modify the values of respective header fields in the packet.
- *Change-TTL:* The various Change-TTL actions modify the values of the IPv4 Time To Live (TTL), IPv6 Hop Limit, or MPLS TTL in the packet.

An *Action Set* is a list of actions associated with a packet that are accumulated while the packet is processed by each table and executed when the packet exits the processing pipeline. Instructions are of four types:

- *Direct packet through pipeline:* The Goto-Table instruction directs the packet to a table farther along in the pipeline. The Meter instruction directs the packet to a specified meter.
- *Perform action on packet:* Actions may be performed on the packet when it is matched to a table entry.
- *Update action set:* Merge specified actions into the current action set for this packet on this flow, or clear all the actions in the action set.
- *Update metadata:* A metadata value can be associated with a packet. It is used to carry information from one table to the next.

Flow-Table Pipeline

A switch includes one or more flow tables. If there is more than one flow table, they are organized as a pipeline as shown in Figure 4, with the tables labeled with increasing numbers starting with 0.

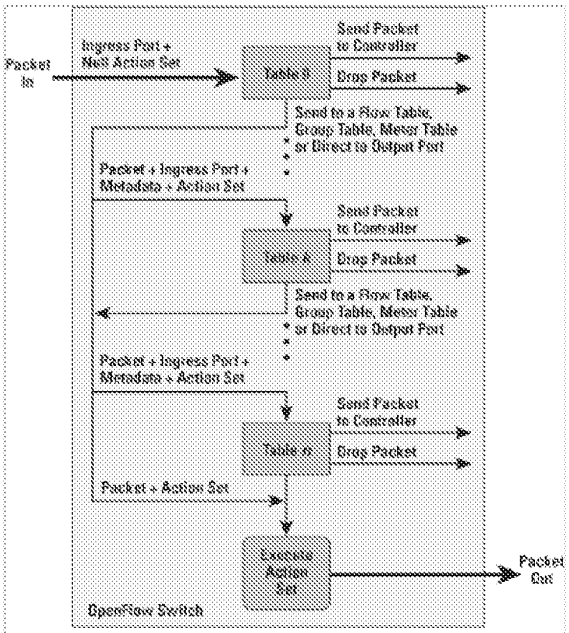


Figure 4: Packet Flow Through OpenFlow Compliant Switch

When a packet is presented to a table for matching, the input consists of the packet, the identity of the ingress port, the associated metadata value, and the associated action set. For Table 0, the metadata value is blank and the action set is null. Processing proceeds as follows:

1. Find the highest-priority matching flow entry. If there is no match on any entry and there is no table-miss entry, then the packet is

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0022

dropped. If there is a match only on a table-miss entry, then that entry specifies one of three actions:

a. Send packet to controller. This action will enable the controller to define a new flow for this and similar packets, or decide to drop the packet.

b. Direct packet to another flow table farther down the pipeline.

c. Drop the packet.

2. If there is a match on one or more entries other than the table-miss entry, then the match is defined to be with the highest-priority matching entry. The following actions may then be performed:

a. Update any counters associated with this entry.

b. Execute any instructions associated with this entry. These instructions may include updating the action set, updating the metadata value, and performing actions.

c. The packet is then forwarded to a flow table further down the pipeline, to the group table, or to the meter table, or it could be directed to an output port.

For the final table in the pipeline, forwarding to another flow table is not an option.

If and when a packet is finally directed to an output port, the accumulated action set is executed and then the packet is queued for output.

OpenFlow Protocol

The OpenFlow protocol describes message exchanges that take place between an OpenFlow controller and an OpenFlow switch. Typically, the protocol is implemented on top of SSL or Transport Layer Security (TLS), providing a secure OpenFlow channel.

The OpenFlow protocol enables the controller to perform add, update, and delete actions to the flow entries in the flow tables. It supports three types of messages, as shown in Table 1.

- *Controller-to-Switch:* These messages are initiated by the controller and, in some cases, require a response from the switch. This class of messages enables the controller to manage the logical state of the switch, including its configuration and details of flow- and group-table entries. Also included in this class is the Packet-out message. This message is used when a switch sends a packet to the controller and the controller decides not to drop the packet but to direct it to a switch output port.
- *Asynchronous:* These types of messages are sent without solicitation from the controller. This class includes various status messages to the controller. Also included is the Packet-in message, which may be used by the switch to send a packet to the controller when there is no flow-table match.
- *Symmetric:* These messages are sent without solicitation from either the controller or the switch. They are simple yet helpful. Hello messages are typically sent back and forth between the controller and switch when the connection is first established. Echo request and reply messages can be used by either the switch or controller to measure the latency or bandwidth of a controller-switch connection or just verify that the device is operating. The Experimenter message is used to stage features to be built into future versions of OpenFlow.

Table 1: OpenFlow Messages

Name	Description
Controller-to-Switch	
Features	Request the capabilities of a switch. Switch responds with a features reply that specifies its capabilities.
Configuration	Set and query configuration parameters. Switch responds with parameter settings.
Modify-State	Add, delete, and modify flow/group entries and set switch port properties.
Read-State	Collect information from switch, such as current configuration, statistics, and capabilities.
Packet-out	Direct packet to a specified port on the switch.
Barrier	Barrier request/reply messages are used by the controller to ensure message dependencies have been met or to receive notifications for completed operations.
Role-Request	Set or query role of the OpenFlow channel. Useful when switch connects to multiple controllers.
Asynchronous-Configuration	Set filter on asynchronous messages or query that filter. Useful when switch connects to multiple controllers.
Asynchronous	
Packet-in	Transfer packet to controller.
Flow-Removed	Inform the controller about the removal of a flow entry from a flow table.
Port-Status	Inform the controller of a change on a port.
Error	Notify controller of error or problem condition.
Symmetric	
Hello	Exchanged between the switch and controller upon connection startup.
Echo	Echo request/reply messages can be sent from either the switch or the controller, and they must return an echo reply.
Experimenter	For additional functions.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0023

The OpenFlow protocol enables the controller to manage the logical structure of a switch, without regard to the details of how the switch implements the OpenFlow logical architecture.

Summary

SDNs, implemented using OpenFlow, provide a powerful, vendor-independent approach to managing complex networks with dynamic demands. The software-defined network can continue to use many of the useful network technologies already in place, such as virtual LANs and an MPLS infrastructure. SDNs and OpenFlow are likely to become commonplace in large carrier networks, cloud infrastructures, and other networks that support the use of big data.

References

[1] Greg Goth, "Software-Defined Networking Could Shake Up More than Packets," *IEEE Internet Computing,* July/August, 2011.

[2] Robin Layland, "The Dark Side of Server Virtualization," *Network World,* July 7, 2010.

[3] Open Networking Foundation, "Software-Defined Networking: The New Norm for Networks," ONF White Paper, April 12, 2012.

[4] Dell, Inc., "Software Defined Networking: A Dell Point of View," Dell White Paper, October 2012.

[5] Steven Vaughan-Nichols, "OpenFlow: The Next Generation of the Network?" *Computer,* August 2011.

[6] Thomas A. Limoncelli, "OpenFlow: A Radical New Idea in Networking," *Communications of the ACM,* August 2012.

[7] Open Networking Foundation, "OpenFlow Switch Specification Version 1.3.0," June 25, 2012.

[8] Sean Michael Kerner, "OpenFlow Protocol 1.3.0 Approved," *Enterprise Networking Planet*, May 17, 2012.

WILLIAM STALLINGS is an independent consultant and author of many books on security, computer networking, and computer architecture. His latest book is *Data and Computer Communications* (Pearson, 2013). He maintains a computer science resource site for computer science students and professionals at ComputerScienceStudent.com. He has a Ph.D. in computer science from M.I.T. He can be reached at ws@shore.net

Information For
Small Business
Midsize Business
Service Provider
Industries
Marketplace
Contacts
Contact Cisco
Meet our Partners
Find a Reseller

News & Alerts
Newsroom
Blogs
Field Notices
Security Advisories
Technology Trends
Cloud
Internet of Things (IoT)
Software Defined Networking (SDN)

Support
Downloads
Documentation
Communities
DevNet
Learning Network
Support Community
Video Portal

About Cisco
Investor Relations
Corporate Social Responsibility
Environmental Sustainability
Trust and Transparency Center
There's Never Been A Better Time
Careers
Search Jobs
We Are Cisco
Programs
Cisco Designated VIP Program
Cisco Powered
Financing Options

Contacts | Feedback | Help | Site Map | Terms & Conditions | Privacy Statement | Cookie Policy | Trademarks

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0024

OpenRoads: Empowering Research in Mobile Networks

Kok-Kiong Yap, Masayoshi Kobayashi, Rob Sherwood, Nikhil Handigol,
Te-Yuan Huang, Michael Chan, and Nick McKeown
Stanford University, NEC System Platforms Research Labs, and Deutsche Telekom Inc., R&D Lab
yapkke@stanford.edu, m-kobayashi@eo.jp.nec.com, robert.sherwood@telekom.com,
{nikhilh,huangty,mcfchan,nickm}@stanford.edu

1. INTRODUCTION

Many researchers are working on mobility management. Continued innovation and development in this area can naturally be expected. We anticipate rapid progress will be made in the areas of handover between different technologies, separation of "virtual" wireless service providers from the physical infrastructure, lossless handovers and multi-homing solutions. There is no lack of good ideas. However, testing and verifying these ideas at a significant scale has proven to be difficult. The community will therefore benefit from a common platform, in which they can reuse and build upon each other's work. To this end, we developed *OpenRoads*, a platform for innovation in mobile networks.

The goal of OpenRoads is to create an *open* platform in which other researchers and ourselves can experiment and explore different mobility solutions, network controllers and/or routing protocols. OpenRoads provides researchers control of the network through two means, namely, control of the datapath using OpenFlow [2, 3] and control of the device configuration using SNMP.

Via control given by OpenFlow and SNMP, we further provide abstract events and functions for network management in OpenRoads. The abstraction in turns allows us to manage with ease heterogeneous network devices, namely different wireless technologies such as WiFi and WiMAX. Further, sharing the vision of a virtualized network, such as in [1, 5], OpenRoads also presents our attempt to share network control among different researchers.

Through our poster, shown in Fig. 1, we present OpenRoads to the community at large. Our long term goal is to develop OpenRoads into a robust and effective platform for research in mobile networks, that is adopted by other campuses, for experimentation in their production networks.

2. CONTROLLING A NETWORK SLICE

A major distinction between today's cellular network and IP-based network is the control plane. IP has the control plane embedded in each and every router, resulting in an inflexible infrastructure with glacially slow innovation. On the other hand, today's cellular network separates the control plane and rapid innovations in the control plane are common.

A typical Ethernet switch consists of the flow table in hardware and the control logic in software, all of which is enclosed in a single device. Adopting a similar philosophy as the cellular network, OpenFlow separates the control logic of a switch from the physical device, exporting it to an external controller. By controlling the flow table of a switch from an external controller, OpenFlow allows rapid innovation to the datapath of the network. To deploy OpenFlow in our closet (and closets of other researchers), we are adding OpenFlow as feature to commercial switches and wireless devices, such as WiFi APs and WiMAX base-station.

Building on OpenFlow, we developed a platform OpenRoads, which allows researchers to control the network via OpenFlow and SNMP. Our resulting architecture is as shown in right half of Fig. 1. While OpenFlow provides control of forwarding, SNMP provides control of device configuration. Being a widely deployed interface, SNMP allows researchers to configure the switches and wireless access points in a network. Configuration is especially critical in wireless devices, where parameters like transmit power can directly impact on the datapath performance. Further, controllers can also be promptly informed of network events such as a host joining a WiFi AP.

To further share this control of the network among multiple researchers, we use the FlowVisor [4] and SNMP demultiplexer to flexibly divide the control. This is achieved through appropriate proxying, rewriting and policing of the control messages. More specifically, each experiment is given a "flow-space" (a range of header values). By ensuring that experiments cannot exert control beyond their designated "flow-space", FlowVisor provides isolation between experiments. Also, by appearing as a controller to network devices and as devices to the controllers, FlowVisor (and SNMP demultiplexer) allows researchers to develop as if they have complete over the entire network. This allows the scale of the experiment to be expanded without any change in code.

Finally, to allow for rapid development on OpenRoads, we provide abstract high-level control interfaces, such as redirection of a flow in the network and retrieving all flows associated with a host. These interfaces can be categorized as communication (i.e., messaging between devices), flow management (e.g., install flow entries for specific route), and device control (e.g., set power level of WiFi APs).

3. PROGRESS

OpenFlow has been ported as a feature in many switches from multiple vendors, including NEC and HP. Many of these devices already supports SNMP. We have also extended OpenFlow and SNMP into WiFi APs built with PCEngine boxes running Linux, and a NEC WiMAX basestation. Armed with these devices, we have deployed a wireless access network consisting of 5 switches, 30 WiFi APs and a WIMAX basestation. With this physical infrastructure, we are able to validate the usefulness of OpenRoads.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0025

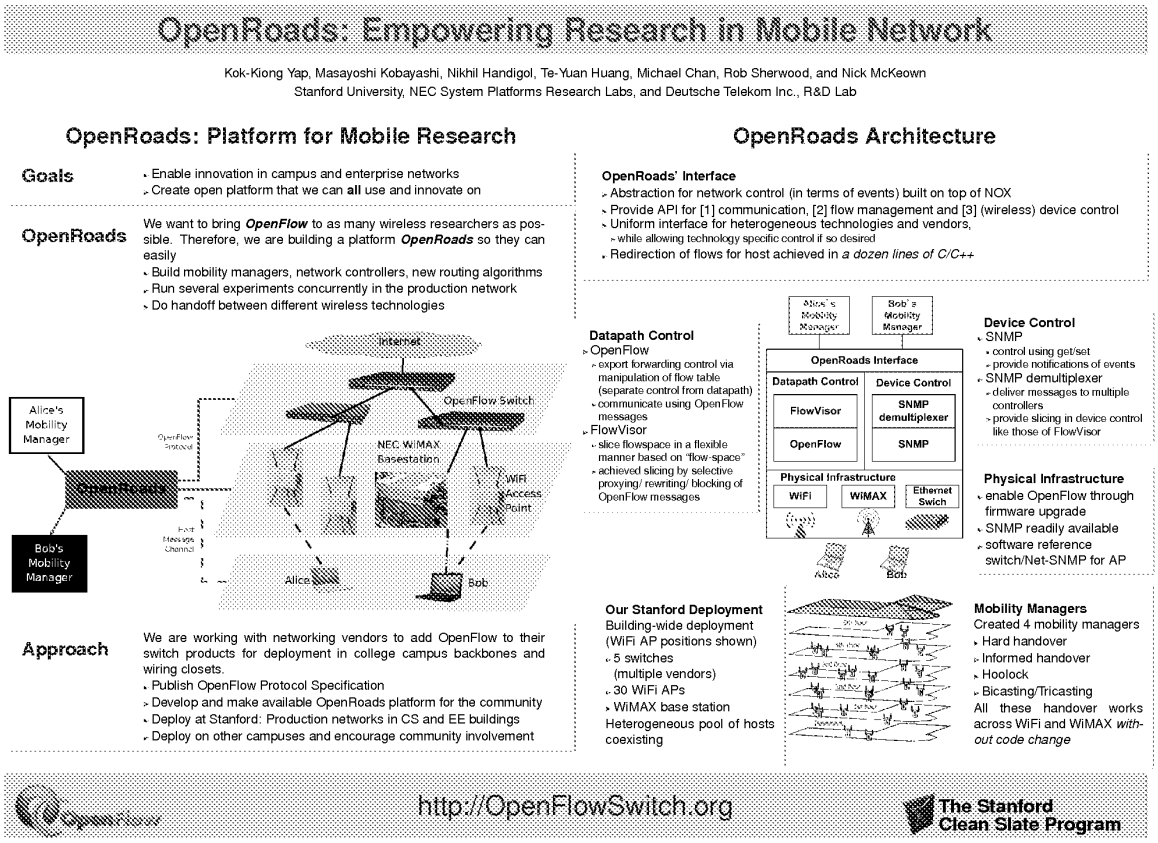


Figure 1: OpenRoads' Poster

Using our OpenRoads platform, we were able to create multiple mobility managers (some of which were created as part of a class) and run them concurrently on a single physical infrastructure. The mobility managers includes hard handover, informed handover, bicasting/tricasting and Hoolock. In informed handover, the client is provided with information to aid handover decisions. While in bicasting/tricasting, flows are multicasted to the client over multiple interfaces. Hoolock is a novel handover scheme which also exploits multiple interfaces on a device to achieve lossless handover, without the need to reorder at the client. All of the above mobility managers are able to perform handover between WiFi and WiMAX without any code change. In these mobility managers, OpenRoads' interface allows us to redirect flows associated with a host in 12 lines of C/C++.

While running these mobility managers, we have a heterogeneous pool of hosts coexisting in the network. For example, the bicasting mobility assumes all devices have more than one network interfaces, while hard handover assumes the contrary, i.e., all devices has one interface. These devices coexists in our network, and are handled by their respective mobility managers without conflict.

We have created a network in which multiple experiments can be run concurrently over the same physical infrastructure. In the long run, we hope OpenRoads will inspire the creation of virtual service providers that are independent of the infrastructure, allowing for seamless handover between different wireless technologies.

4. REFERENCES

[1] N. Feamster, L. Gao, and J. Rexford. How to lease the internet in your spare time. *SIGCOMM Comput. Commun. Rev.*, 37(1):61–64, 2007.

[2] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, April 2008.

[3] The OpenFlow Switch Consortium. http://www.openflowswitch.org.

[4] Rob Sherwood, Michael Chan,Glen Gibb, Nikhil Handigol,Te-Yuan Huang,Peyman Kazemian,Masayoshi Kobayashi, David Underhill, Kok-Kiong Yap, and Nick McKeown. Carving research slices out of your production networks with OpenFlow. In *Proceedings of ACM SIGCOMM (Demo)*, Barcelona, Spain, August 2009.

[5] A virtual network infrastructure. http://www.vini-veritas.net.

OpenRadio: A Programmable Wireless Dataplane

Manu Bansal, Jeffrey Mehlman, Sachin Katti, Philip Levis
Stanford University
{manub, jmehlman, skatti}@stanford.edu, pal@cs.stanford.edu

Abstract

We present OpenRadio, a novel design for a programmable wireless dataplane that provides modular and declarative programming interfaces across the entire wireless stack. Our key conceptual contribution is a principled refactoring of wireless protocols into processing and decision planes. The processing plane includes directed graphs of algorithmic actions (eg. 54Mbps OFDM WiFi or special encoding for video). The decision plane contains the logic which dictates which directed graph is used for a particular packet (eg. picking between data and video graphs). The decoupling provides a declarative interface to program the platform while hiding all underlying complexity of execution. An operator only expresses decision plane rules and corresponding processing plane action graphs to assemble a protocol. The scoped interface allows us to build a dataplane that arguably provides the right tradeoff between performance and flexibility. Our current system is capable of realizing modern wireless protocols (WiFi, LTE) on off-the-shelf DSP chips while providing flexibility to modify the PHY and MAC layers to implement protocol optimizations.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

General Terms

Design, Management, Performance

Keywords

Programmable, Infrastructure, Dataplane, Wireless

1. INTRODUCTION

Wireless protocols evolve continuously. For example, 3GPP (the standards body responsible for specifying the LTE cellular standard) expects to release major updates to the basic LTE standard once every 18 months, and minor updates continuously. Even within a deployed operational network, operators and vendors need to continuously optimize the network to handle problems such as inter-cell interference, new traffic classes such as video and so on. For example, operators need to implement management mechanisms that dynamically adjust spectrum and power allocation at neighboring basestations at a fine-grained subcarrier level on timescales of hundreds of milliseconds to ensure that mobile handsets at the edge of both cell-sites are not adversely affected. Similarly, to cope with the growing popularity of video over spectrum starved cellular networks, operators would prefer to use a PHY layer optimization such as unequal error protection (UEP) - where certain frames (I-frames) are better protected than others (P- or B-frames) [15][25] - that can deliver much higher video quality while using lesser spectral resources. Similar protocol changes and optimizations abound (some more examples are discussed in Sec. 2) that are beneficial to implement for a variety of reasons.

The wireless infrastructure is expected to support such evolvability. In the past, such changes typically required replacing the basestations, since the protocol definition was closely coupled with the hardware in the form of specific ASICs designed for each protocol. Such an approach worked fine when protocol upgrades were relatively slow (several years) and basestations were deployed sparsely and covered large areas. However, in current networks, protocol changes are continuous (once every few months) and networks are dense (e.g. one basestation per block in urban areas) to provide high capacity and uniform coverage. Physically and frequently upgrading the basestations for such dense networks is prohibitively expensive.

In this paper, we argue that wireless infrastructure needs a programmable dataplane to support such evolvability. Specifically, the basestations should be remotely programmable to enable vendors and operators to upgrade and optimize the network completely in software. To build such functionality, we present OpenRadio, a novel design for a *programmable wireless network dataplane*. Built around a commodity multi-core hardware platform, the core component of OpenRadio is a software abstraction layer that exposes a modular and declarative interface to program the PHY (baseband) and MAC layers. An operator describes a protocol by defining *rules* to match subsets of traffic streams and specifying *actions* to process them. Rules are logical predicates on parameters of packets such as header fields, received signal strength, channel frequency and other fields that may be programmed. Actions describe behavior such as encoding/decoding of data and scheduling of traffic on the channel.

OpenRadio's design makes two contributions. First, it decouples wireless protocol definition from the hardware, even while ensuring that commodity multi-core platforms can be used for implementing the protocols. Such decoupling is feasible because we can identify low-level atomic building blocks in wireless protocols at the PHY

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
HotSDN'12, August 13, 2012, Helsinki, Finland.
Copyright 2012 ACM 978-1-4503-1477-0/12/08 ...$15.00.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0027

and MAC layers that are shared across different protocols (different versions of LTE 4G, 3G, WiFi and so on). Further, OpenRadio shows that commodity multi-core platforms are capable of executing such atomic blocks efficiently. Consequently, defining and customizing protocol definition programmatically is feasible, as long as modular abstractions exist to re-program individual protocol elements.

OpenRadio's second contribution is the design of this software abstraction layer that exposes a modular and declarative interface to program wireless protocols. The key idea is a principled decomposition of wireless protocols into separate *processing* and *decision* plane components with a simple, well-understood API between them. This decomposition is the main enabler of an efficient implementation of a flexible and modular wireless programming interface. The processing plane includes *actions*: operations that only involve data-manipulation on packet data but no branching logic. The decision plane includes *rules*: all decision logic that picks sequences of operations to transform data stream subsets. Such decomposition lends itself to efficient real-time implementation on commodity digital signal processors (DSPs) that are widely available from a number of chip vendors [26, 12, 9, 14].

OpenRadio can be used to specify both the underlying protocols as well as optimizations. For example, a rule for WiFi would require all received packets matching the 6Mbps data rate in the PHY header to be processed with the action of BPSK demodulation. Protocol optimizations can be expressed by adding rules and actions to an existing definition. For video delivery, an operator using OpenRadio might optimize for this class of traffic by inserting a new rule to match all video packets (based on a header field). The rule would require them to be processed through a special action of UEP (en)decoding. The UEP (en)decoder action would also be programmed by the operator as an atomic block. Meanwhile, the existing standard WiFi-data handling remains untouched.

More generally, our longer-term goal is to design a programmable wireless dataplane that includes the basestations, backhaul network switches/routers and gateways to realize a software defined cellular infrastructure. Our goal in this paper, however, is more modest. It is to discuss the design of a programmable basestation that forms one (albeit important) component of this infrastructure. At this stage, we have a prototype implementation of OpenRadio that supports the WiFi PHY on off-the-shelf multicore DSP processors from TI [26]. The platform is capable of fully processing the WiFi PHY protocol in software running at the highest speeds (upto 108Mbps) and meeting all protocol deadlines. We pick WiFi for proof-of-concept since its processing requirements are the most stringent due to wide-bandwidth, multi-carrier modulation and random-access deadlines on the order of microseconds. We plan to add support for LTE in the near future, as well as leverage the dataplane to start researching the design of a control plane for cellular SDN infrastructure.

In the following section, we discuss more specific examples of the kinds of protocol changes that operators would want to implement in their networks remotely, and that OpenRadio's design supports.

2. USE CASES

Programmable basestations for cellular infrastructure are being increasingly favored over fixed-function hardware as exemplified by offerings of "software-upgradable" platforms for HSPA/WCDMA/LTE from some of the top vendors [23, 3]. Still, many desirable features either remain hard to program or are simply not possible to implement. A well-designed programmable dataplane can enable these features. Let us consider some specific use-cases to gain insight into the kind of programmability required.

2.1 Cell-size based optimization

OFDMA leads to power-amplifier (PA) inefficiencies due to a high peak-to-average-power ratio (PAPR) as compared to single-carrier modulation. This lead to the choice of SC-FDMA for LTE (Rel 8/9) uplink to save user equipment (UE) power. However, advances in PA technology have now led to the adoption of OFDMA (Rel 10) in favor of higher spectral efficiency. This is especially true for small-cells where the UE needs to transmit at a lower power than a macro-cell. The switch from SC-FDMA to OFDMA effectively involves toggling a DFT block on the PHY signal chain. Definitions of individual processing blocks remain unchanged.

Another cell-size based optimization could involve picking the right FFT size to balance multipath effects against PAPR inefficiencies. A higher FFT size leads to better multipath resiliency but a higher PAPR value which is less energy efficient. A smaller cell can use a smaller FFT size for the same data delivery performance since it suffers less multipath fading and delay spread. Similar FFT size adaptation based on link length can also benefit WiFi. The FFT processing block only needs to be configurable in FFT size without any changes in the algorithm.

2.2 Co-existence of heterogeneous cells

Cellular networks are rapidly becoming heterogeneous in cell size with deployment of small-cells. Femto-cells enhance coverage in a small area such as a home but generally provide access to a closed group of subscribers. However, deployment in the same frequency as the containing macro-cell can cause strong interference to external users in the same region. This has led to techniques like Almost Blank Subframes (ABS) [6] that co-ordinate the use of time-frequency resource elements between the two cells to avoid interference. Such co-ordination can be implemented with programmability in the MAC scheduler that decides whether or not to use a resource element at any given time.

2.3 Application-specific wireless service

A promising feature of LTE is multiple service profiles [4] with varying throughput and latency characteristics. The intent of the QoS profiles is to provide customized network service to diverse applications. However, the standard only specifies profile parameters while leaving the mechanisms open to implementation. A programmable dataplane provides the vehicle necessary to implement various profiles and allowing their extension to new application classes. Different processing chains can be tied to different traffic classes. Prioritization, scheduling of traffic and error recovery mechanisms can also be programmed according to application needs. Those features can also evolve with new applications such as public safety networks, urban sensing and media broadcast in addition to classic voice, video and data delivery. A programmable dataplane can similarly facilitate implementation of QoS features for WiFi, WiMAX, White-space networking etc.

2.4 Evolving standards

Cellular wireless standards have been upgrading every couple of years [1, 2]. Users are ready to consume network upgrades at an even faster pace with an average UE lifespan of 18 months [10]. Yet, standards take a decade or more to reach their peak market share [11]. Arguably, the biggest impediment to evolution of the infrastructure is its inability to be programmed in-place and incrementally.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0028

Wireless networks are already operating with very high spectral efficiency (within $\approx 20\%$ of capacity limits). Consequently, signal processing motifs are stabilizing with little room for improvement. Moreover, cellular networks are designed such that most intelligence is place in the infrastructure while keeping UEs simple. This allows a bigger playing fields on the base-station side that is still interoperable with existing users. Most gains are to be had by optimizing networks according to the operational scenario and load patterns, which tend to be very diverse [11]. At the same time, the set of users is always a rich mix of capabilities as new standards take long to penetrate. The ideal base-station must be able to cater to all of them simultaneously. OpenRadio provides a platform where new protocols and features can be "plugged-in" seamlessly. Its programmable dataplane provides flexibility in selection, configuration and arrangement of data processing motifs to compose protocols optimized to the operational scenario.

3. DESIGN GOALS

Our goal is to design a programmable wireless dataplane with a modular interface that can provide both the performance and flexibility necessary for implementing a wide class of modern wireless protocols and their optimizations.

As a *performance* benchmark, the platform must be capable of processing a bandwidth of at least 20MHz for OFDM-complexity protocols, including but not limited to IEEE 802.11 WiFi, IEEE 802.16 WiMAX and 3GPP LTE-Advanced. The platform must provide enough computation throughput for these applications as well as the processing latency requirements associated with specific protocols. In the case of WiFi, the most challenging deadline is an ACK-turnaround time of $25\mu s$ for every successfully decoded data packet.

Flexibility is a more qualitative characteristic than performance. Our design goal is to provision the system to allow a smooth tradeoff between performance and deviation from the standard configuration. Small variations in behavior should be enabled with little or no performance loss. For example, a version of WiFi optimized for video delivery using UEP only requires a change in the encoding of I- and B-/P-frame bits into a packet of the PHY layer. All other operations are preserved. Consequently, the platform must be able to deliver near-WiFi computation performance for video-WiFi. At the same time, the platform should not preclude more significant changes from being realized at all.

An equally important design goal is making the OpenRadio platform easy to program. The system must provide *modularity* such that the involvement of a programmer is limited to only the affected components of the protocol. Adapting a WiFi implementation to video-WiFi should not require the operator to modify or even understand the OFDM (de)modulator or the CRC implementation. Without such modularity, the utility of a flexible dataplane will be far from practical.

4. PROGRAMMING ABSTRACTIONS

The OpenRadio dataplane design provides a network substrate that can be programmed using a modular, declarative interface, provides high performance capable of supporting the highest rate modern wireless protocols, and can be built using commodity DSP silicon. In this section, we describe its declarative interface.

4.1 Modular declarative interface

OpenRadio's dataplane enables the network operator to declaratively program the network stack for specific traffic subsets. He only specifies what packet processing behavior he expects from the network for different traffic classes. However, the operator does not need to specify how the packet processing behavior is integrated into an operational network and realized in practice. This low-level complexity is abstracted out by the interface.

OpenRadio's declarative interface is built on a *principled refactoring of decision and processing paths in wireless protocols*. Specifically, we observe that protocols can be finely partitioned into two separate parts: processing blocks which specify how a protocol transforms analog waveforms into bits, and decision logic which specifies when different processing transformations are used. For example, the 54Mbps OFDM decoding sequence of 802.11g is a processing specification, while interpreting the signal field of the 802.11 PHY layer control header to determine that the payload will be at 54Mbps is a decision specification.

The processing plane encapsulates the series of signal processing algorithms used in a PHY processing chain where each module corresponds to a single algorithm. For example, the processing plane will have modules for FFT, 64-QAM mapping and slicing, convolutional encoding, Viterbi decoding, etc. These modules correspond to *blocks* in a signal processing chain. The decision plane encapsulates all the decision logic functionality. It is responsible for selecting processing plane blocks to be used for encoding or decoding a packet. For example, the decision plane logic for 802.11g will appropriately configure and stitch the processing plane plane blocks to decode a 54Mbps packet without needing to touch any block internals. Similarly, a protocol modification that rearranges the same signal processing blocks would only need to modify the decision plane, leaving the processing plane blocks untouched. As another example, LTE changes the way packets are processed on the downlink and uplink. It uses standard multicarrier OFDM for the downlink, while it uses a modified single carrier modulation (SC-FDM) for uplink to conserve mobile client batteries [5]. The details of the signal processing algorithms are unimportant, but the main take-away is that both OFDM and SC-FDM are simply different ways of configuring and stitching basic signal processing blocks such as FFTs, channel decoders and slicers. The decision plane can simply specify such re-configurations and re-arrangements to realize different protocol optimizations.

The two key choices - granularity of processing plane blocks and decoupling decision plane logic from the processing plane - are directly motivated by the goal of modular programmability. Decoupling allows logical changes and rearrangement of signal processing blocks to be limited to the decision plane without requiring any understanding of the processing plane. Reconfiguring the behavior of a processing plane block limits changes to the specific block without affecting other blocks or the decision plane. Adding a new block (eg. a new channel encoder or decoder optimized for video) is similarly a contained effort. The level of abstraction of processing plane blocks (eg. FFT vs. low level instructions such as complex multiply and accumulate) is high enough to allow most programmability to be expressed declaratively and allows convenient reuse of code. The imperative part is abstracted out in processing plane blocks. Such separation of processing and decision planes allows an intuitive programming model made up of rules and actions.

4.2 Rule-action programming model

The decision plane of a protocol has a natural representation as a set of *rules* encoding the decision logic. These rules govern the processing plane by naming *actions* that must be performed depending on the outcome of rule evaluation. For example, a rule-action pair will read `if hdr.rate == 54Mbps, perform action_decode_54Mbps`.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0029

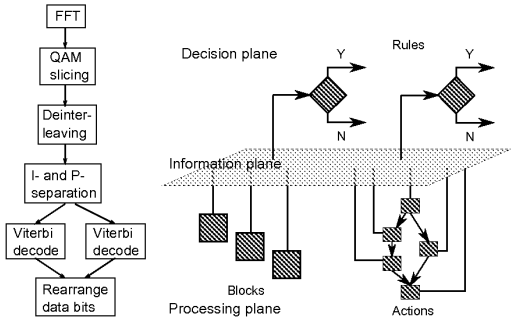


Figure 1: (a) Directed acyclic graph of an action customized for video delivery. (b) OpenRadio programming model depicting the separation of decision and processing planes and their relation with the information plane.

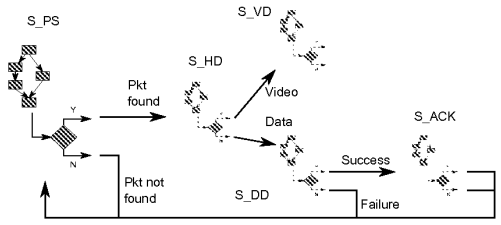


Figure 2: A state machine with states for packet search (S_PS), header decode (S_HD), video decode (S_VD), data decode (S_DD) and sending an acknowledgement (S_ACK).

Actions are directed acyclic graphs (DAGs) of signal-processing operations. Each operation is typically simple: it takes values from an input stream, applies a simple operation such as an FFT, and outputs the processed stream. Operations do not include conditionals or other complex, dynamic logic since that is factored out into the decision plane. An example action DAG for a video decoding chain is shown in fig. 1(a).

The regular interface of a processing plane block is composed of the data inputs (eg. complex samples) and data output (eg. soft bits, complex samples). In addition, a block has configuration inputs (eg. FFT size) and output statistics (eg. energy value for a packet-detect block, rate field for a header-parsing block). Collection of all configurations and statistics from the processing plane forms an *information plane* that acts as the bi-directional interface between decision and processing planes (fig. 1(b)). Rules are written on elements of the information plane. Similarly, the decision plane exercises control over the processing plane through configuration primitives made available on the information plane.

The programming model of rules and actions is a specific concrete representation of the concept of decision and processing plane decoupling. It lets us express protocols declaratively by preserving the decoupling semantics. Actions only specify the functional behavior and data dependencies without specifying any execution schedules, hardware placement details or resource allocation. Similarly, rules only specify the logical content without dictating how they will be implemented. Thus, the model provides a generic API without being tied to any particular system design or hardware platform architecture. It also allows us to leverage extensive prior work where the stream computation in actions is naturally expressed using dataflow languages [16]. The current implementation of OpenRadio uses annotated C code, but its decoupled programming model makes it simple to integrate new domain-specific dataflow programming languages.

4.3 State machine model and deadlines

Rules and actions are assembled into a state machine to express a fully-functional protocol, as in fig. 2. Each state specifies exactly one action to be performed and contains rules for transitioning to other states. For example, the header decode state will perform the action of decoding the header field and include rules to transition to data-decoding states, one for each data rate. Such a state machine model matches naturally with the specification of protocols, some of which use formal state machine languages in the standard itself.

In addition to rules and actions, a first class component of wireless protocols is deadlines. Deadlines may apply to specific actions such as a maximum of $1\mu s$ delay in assessing channel idle/busy state for effective channel backoff implementation in WiFi (i.e. slot-time stipulation), or to a series of actions such as a maximum delay of $25\mu s$ in decoding a packet from the end-of-arrival (due to WiFi ACK-turnaround stipulation.) In the latter case, they exist between certain states, and consist of hard timing requirements on all of the computation (decision and processing-plane) that occurs in the interval between those two states. Valid deadlines may exist between any two states in a protocol but there must be a logical route through the rules from the initial state to the end state.

Similar to the choice of rule-action model, the state machine model is a specific choice for representing protocols that is intuitive and preserves the declarative nature of the programming interface including specification of deadlines. Further, it allows us to leverage prior work on formal state machine languages like [7] with strong semantics to facilitate static analysis that is critical to providing guaranteed execution times and meeting deadlines as we discuss further in sec. 5.

4.4 Challenges

The key challenge in OpenRadio is efficient realization of the decoupled declarative programming model. In particular, the system must be able to meet protocol deadlines reliably without requiring the programmer to optimize implementation specifics. We can summarize the challenges as follows:

- Efficient execution of computationally demanding processing plane actions ($\approx$100 Gflops) while retaining hardware abstraction, programming modularity and flexibility.
- Maintaining realtime property over a sequence of states, i.e., under serial composition, in order to meet deadlines reliably.

The rich programming interface made up of rules-actions and state-machine models is the key feature that distinguishes OpenRadio from a bare-bones software-radio platform or a programmable substrate like an FPGA platform that may offer similar performance. Let us discuss the principles behind an efficient implementation of this programming interface.

5. DESIGN AND IMPLEMENTATION

Decoupling of decision and processing planes allows us to provide a modular and declarative programming interface through the rule, action and state machine models. Can we design a system that implements such an interface with the level of computational performance needed for modern protocols with 20-80MHz of bandwidth while meeting tight latency deadlines on the order of tens of microseconds? We answer the question in the affirmative in this

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0030

Processing plane block	WiFi	LTE	3G	DVB-T
FIR/IIR	✓	✓	✓	✓
Correlation	✓	✓	✓	✓
Spreading			✓	
FFT	✓	✓		✓
Channel Estimation	✓	✓	✓	✓
QAM Mapping	✓	✓	✓	✓
Interleaving	✓	✓	✓	✓
Convolutional Coding	✓	✓	✓	✓
Turbo Coding		✓	✓	
Randomization	✓	✓	✓	✓
CRC	✓	✓	✓	

Table 1: Many computationally intensive processing plane blocks are common to all modern PHYs, differing only in configuration.

section describing the design and early results demonstrating its feasibility. Our system is based on two keys principles.

5.1 Design principles

Learning from existing programmable solutions, OpenRadio follows *the principle of scoping flexibility*. We aim to provide enough flexibility to implement most protocols and optimizations of interest but no more. For example, the FFT block only needs to provide configurable FFT-length, not a choice of different Fourier Transform algorithms. Similarly, the convolutional decoder needs to be configurable in the coding rate (say, 1/2, 1/3, 1/4, 1/5) or constraint length over small parameter ranges but the actual algorithm (Viterbi or BCJR) is immaterial. There is, however, still scope for blocks that need almost arbitrary flexibility such as a bit interleaver.

Restricting the scope of flexibility where advantageous allows for optimizing the hardware to perform some of the relatively-fixed operations extremely efficiently while being able to provide the requisite configurability. The merit of this approach lies in the observation that many of the heavy computational processing plane blocks of modern PHY protocols are common to most standards and fall in this category (Table 1). Further, commodity DSP platforms with hybrid architectures optimizing common *compute kernels* are becoming widely available, either through loosely-coupled on-chip accelerators as in TI KeyStone [26] or through tightly-coupled special-purpose instructions as in CEVA-XC Family [9], [20][14][12]. A well-designed accelerator can provide speed-ups of up to an order of magnitude (10-25x) over pure-software [13]. As a price reference point, a TI KeyStone chip housing tens of accelerators is marked at less than $200 and is capable of processing 80MHz WiFi-like intensity of computation, i.e. it easily supports WiFi, WiMAX, LTE and W-CDMA.

While the above hardware substrate provides the raw horsepower needed for a programmable dataplane, OpenRadio's dataplane must provide a modular declarative interface while ensuring that packets are processed within the deadlines of the protocol. To this end, OpenRadio follows *the key principle of decision and processing plane separation throughout its system design*. In contrast with providing modularity and declarative interfaces, the motivation here is to achieve near-deterministic execution times for processing plane actions. On a hybrid architecture, blocks implemented through hardware accelerators would finish in predictable amounts of time by default. This may not be true of other software-based blocks that will run on CPUs or DSPs. The algorithmic reason for such unpredictability is presence of branches in arbitrary code. For example, a soft-slicer block that could do any of BPSK, QPSK, 16-QAM or 64-QAM, will run for different amounts of time in each mode. However, a branch is precisely a decision plane rule. The factoring of algorithm and code into decision and processing planes pushes out all such branches to the decision plane, leaving only branch-free code in processing plane blocks.

Combining the decoupling principle with assisting implementation choices, we not only achieve deterministic-runtime blocks but also preserve the property under composition as required for actions. This is the key property that allows OpenRadio to systematically meet execution deadlines. An added advantage of this principle is a direct mapping of declarative programs to execution primitives which greatly simplifies the system design.

5.2 Preliminary implementation

The design of OpenRadio system is focused towards utilizing multi-core DSP architectures. They typically contain DSP cores optimized for signal processing computation and hardware accelerators providing speed-up of specific, commonly used algorithms. Such platforms provide the most desired tradeoff between performance and flexibility. Our specific hardware choice is the TI KeyStone architecture [26]. The software challenge is to harness the raw compute horsepower while retaining modular and declarative abstractions.

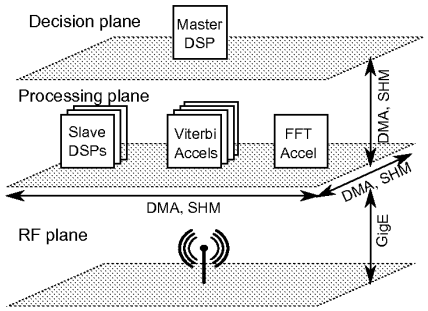


Figure 3: OpenRadio basestation internals with different kinds of compute cores and an RF frontend.

The core objective of the runtime system is *fast and deterministic execution of actions*. While efficient execution is achieved through accelerated hardware, determinism in runtime is obtained through decision/processing separation down to the software level, mimicking the programming model itself. The system architecture is shown in fig. 3. The same separation is applied to hardware resources as well, designating the decision plane core as the master core controlling the processing plane slave cores and accelerators.

The decision plane runtime emulates the protocol state-machine by evaluating rules as transitions are made between states. The action associated with each state is computed on the processing plane by assigning constituent blocks to slave cores (or accelerators) according to a pre-computed execution and resource allocation schedule. Such static scheduling is made possible because of our specific implementation choices such as software-addressed L2 memory and event-polling instead of interrupts. They eliminate major sources of indeterminism from the underlying hardware architecture as well. Further, using a message-passing model for inter-processor communication allows accurate accounting of IPC latencies. Full description of implementation and scheduling framework are beyond the scope of this paper. However, it is worth stating that the scheduling of actions in our deterministic model bears strong resemblance with instruction loop-pipelining in compilers. This reduction of model allows us to leverage compiler scheduling algorithms as in [24].

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0031

Our early prototype is capable of processing full 54Mbps 802.11g PHY layer stack on a single KeyStone chip with more than one full DSP core worth of computer power still available for implementing protocol modifications. The utilization figure includes the overhead introduced by the runtime software that strives to maintain modularity. This overhead measures to roughly 20% of the entire compute load.

5.3 Design limitations

Our software system design applies equally well to other multi-core processors, hybrid or symmetric. Still, it is based on certain judicious assumptions.

First, OpenRadio assumes that the majority of the computation time is spent in the processing plane of the protocol, and decision overhead is negligible. However, heavy decision planes would incur prohibitive inter-core communication overheads implying that decision/processing plane separation is not the best design choice. Similarly, a heavy sequential algorithm on the processing plane would restrict the benefits of a multi-core architecture if it forbids pipelined parallelism.

Second, OpenRadio assumes decision/processing refactoring is meaningful in signal processing blocks. However, some blocks that need to take decisions, like packet-search block, are best expressed with some embedded decision code. If the gap in average and worst-case run times becomes too high, maintaining the deterministic semantic on that block could become very inefficient.

6. RELATED WORK AND DISCUSSION

OpenRadio has a lot in common with SDN efforts such as OpenFlow [18], both in goals and in challenges. The key design consideration is to pick the right abstractions to balance flexibility against performance. However, there are some major differences too. Interoperation with existing devices on the Internet is a major consideration for layer-3 SDN interfaces. For OpenRadio, inter-operation is much less of an issue as it operates on layers 1 and 2 which are contained within autonomous domains. Further, a programmable wireless dataplane has a massive computational load that must be processed under strict deadlines, a challenge unique to wireless.

There are comparable solutions for programmable wireless dataplanes that excel in some dimension at the cost of others. GNURadio [8] offers unlimited flexibility through purely software-based processing and convenient programming through high-level interfaces at the cost of performance and real-time guarantees. SORA [27] offers full flexibility and sufficient performance while arguably sacrificing ease of programming by nearly by-passing the desktop operating system. WARP [17] and AirBlue [22] use an FPGA approach to retain flexibility equal to software while gaining performance but end up driving the price-point high. Others like NI PXIe-1082 SDR [21] have taken an all-out approach with unrestrained FPGA and CPU horsepower resulting in extremely expensive systems (≈$50k per SDR platform) only good for small-scale laboratory research and prototyping. Yet other cellular domain platforms [3][26, 12, 9, 14] tend to provide the right price-performance and flexibility point but remain hard to program. This has led to the development of higher-level interfaces [22, 19] to work with these platforms.

OpenRadio strikes a unique combination of the programmability-performance-price metrics by treating all three as primary objectives from the get-go. The design is based on principled refactoring of wireless protocols into processing and decision planes. It is general enough to be implemented on a variety of platforms including DSPs, FPGAs and multi-core desktop machines. Code has enough structure that it can be statically analyzed to yield guaranteed real-time performance and efficient hardware use while still exposing modular, declarative interfaces.

Our intent is to provide programmability across the entire wireless stack. In this paper, most of our design choices are guided by PHY-layer considerations. The decision and processing plane model applies equally well to the MAC layer - the processing plane operation is to map multiple traffic queues onto the channel frequency-time grid while the decision plane operation is to pick the right mapping according to a programmable algorithm. The weighting of computation between the two planes may be different than that in PHY. This exploration is part of our ongoing research.

7. REFERENCES

[1] http://en.wikipedia.org/wiki/3GPP.

[2] http://en.wikipedia.org/wiki/Comparison_of_mobile_phone_standards.

[3] RBS6000 Base Station. *http://www.ericsson.com/ourportfolio/products/base-stations.*

[4] 3GPP TS 23.203 V8.3.1.

[5] 3GPP TS 36.201 - v1.0.0. LTE Physical Layer - General Description.

[6] Alcatel-Lucent. R1-093340: Blank Subframes for LTE. *3GPP TSG RAN WG1,meeting 58, Shenzhen, China, August 2009.*

[7] M. Anlauff. XASM - An Extensible, Component-Based Abstract State Machines Language. In *International Workshop on Abstract State Machines, Lecture Notes on Computer Science (LNCS)*, pages 69–90. Springer-Verlag, 2000.

[8] E. Blossom. GNU radio: tools for exploring the radio frequency spectrum. *Linux J.*, 2004:4–, June 2004.

[9] CEVA, Inc. CEVA-XC High-Performance, Low-Power DSP Cores for Advanced Wireless Communications.

[10] Earth911.com. Facts about cell phones. *http://earth911.com/recycling/electronics/cell-phone/facts-about-cell-phones/.*

[11] Ericsson Corp. Traffic and Market Report, June 2012. *http://www.ericsson.com/res/docs/2012/traffic_and_market_report_june_2012.pdf.*

[12] Freescale, Inc. StarCore Digital Signal Processors.

[13] R. Hameed, W. Qadeer, M. Wachs, O. Azizi, A. Solomatnikov, B. C. Lee, S. Richardson, C. Kozyrakis, and M. Horowitz. Understanding sources of inefficiency in general-purpose chips. ISCA '10. ACM, 2010.

[14] Intel Corporation, Inc. Intel Solutions for the Next Generation Multi-Radio Basestation. *Application Note.*

[15] ISO/IEC Moving Picture Experts Group (MPEG). The MPEG4 Standard. *ISO/IEC 14496 - Coding of audio-visual objects*, 1998.

[16] W. M. Johnston, J. R. P. Hanna, and R. J. Millar. Advances in dataflow programming languages. *ACM Comput. Surv.*, 36(1):1–34, Mar. 2004.

[17] A. Khattab. WARP: A flexible platform for clean-slate wireless medium access protocol design. *SIGMOBILE Mob. Comp. Comm.*, January 2008.

[18] N. McKeown. OpenFlow: Enabling innovation in campus networks. *SIGCOMM Comput. Commun. Rev.*, 38:69–74, March 2008.

[19] Microsoft Research. Brick Specification. *The SORA Project*, 2011.

[20] Mindspeed Technologies, Inc. MindSpeed Baseband Processors.

[21] National Instruments, Inc. NI PXIe-1082 SDR Chassis. *http://sine.ni.com/nips/cds/view/p/lang/en/nid/207346.*

[22] A. Ng, K. E. Fleming, M. Vutukuru, S. Gross, Arvind, and H. Balakrishnan. Airblue: A System for Cross-Layer Wireless Protocol Development. 2010.

[23] Nokia Siemens Networks. Flexi Multiradio BTS. *http://www.nokiasiemensnetworks.com/portfolio/products/mobile-broadband/single-ran-advanced/flexi-multiradio-base-station.*

[24] J. Sánchez and A. González. The Effectiveness of Loop Unrolling for Modulo Scheduling in Clustered VLIW Architectures. ICPP '00. IEEE Computer Society, 2000.

[25] S. Sen, S. Gilani, S. Srinath, S. Schmitt, and S. Banerjee. Design and Implementation of an "Approximate" Communication System for Wireless Media Applications. In *ACM SIGCOMM*, 2009.

[26] Texas Instruments. TMS320TCI6616 Communications Infrastructure KeyStone SoC. *Data Manual.*

[27] J. Zhang. Experimenting software radio with the SORA platform. In *Proceedings of ACM SIGCOMM*, 2010.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0032

2012 European Workshop on Software Defined Networking

Toward Software-Defined Cellular Networks

Li Erran Li
Bell Labs
erranlli@research.bell-labs.com

Z. Morley Mao
University of Michigan
zmao@umich.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

***Abstract*—Existing cellular networks suffer from inflexible and expensive equipment, complex control-plane protocols, and vendor-specific configuration interfaces. In this position paper, we argue that software defined networking (SDN) can simplify the design and management of cellular data networks, while enabling new services. However, supporting many subscribers, frequent mobility, fine-grained measurement and control, and real-time adaptation introduces new scalability challenges that future SDN architectures should address. As a first step, we propose extensions to controller platforms, switches, and base stations to enable controller applications to (i) express high-level policies based on subscriber attributes, rather than addresses and locations, (ii) apply real-time, fine-grained control through local agents on the switches, (iii) perform deep packet inspection and header compression on packets, and (iv) remotely manage shares of base-station resources.**

I. Introduction

The growing popularity of smart phones and tablet computers places an increasing strain on cellular networks. Yet, despite tremendous innovation in mobile applications, the cellular network infrastructure is remarkably brittle. Software defined networking (SDN) can simplify network management, while enabling new services. However, supporting many subscribers, frequent mobility, fine-grained measurement and control, and real-time adaptation introduces scalability challenges that future SDN architectures should address.

A. Today's LTE Cellular Data Networks

Long Term Evolution (LTE) cellular networks [1] connect base stations (eNodeB) to the Internet using IP networking equipment, as shown in Figure 1. The user equipment (UE) connects to a base station, which directs traffic through a serving gateway (S-GW) over a GPRS Tunneling Protocol (GTP) tunnel. The S-GW serves as a local mobility anchor that enables seamless communication when the user moves from one base station to another. The S-GW must handle frequent changes in a user's location, and store a large amount of state since users retain their IP addresses when they move. The S-GW tunnels traffic to the packet data network gateway (P-GW). The P-GW enforces quality-of-service policies and monitors traffic to perform billing. The P-GW also connects to the Internet and other cellular data networks, and acts as a firewall that blocks unwanted traffic. The policies at the P-GW can be very fine-grained, based on whether the user is roaming, properties of the user equipment, usage caps in the service contract, parental controls, and so on.

Besides data-plane functionalities, the base stations, serving gateways, and packet gateways also participate in several

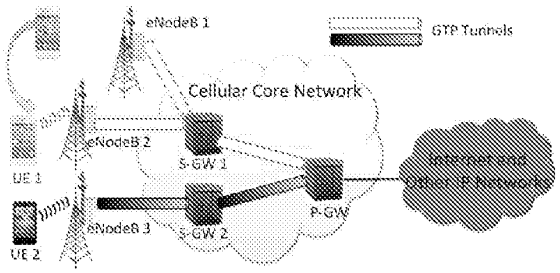


Fig. 1. LTE data plane

control-plane protocols, as illustrated in Figure 2. In coordination with the Mobility Management Entity (MME), they perform hop-by-hop signaling to handle session setup, teardown, and reconfiguration, as well as mobility, e.g., location update, paging, and handoff. For example, in response to a UE's request for dedicated session setup (e.g., for VoIP call), the P-GW sends QoS and other session information (e.g., the TCP/IP 5-tuple) to the S-GW. The S-GW in turn forwards the information to the MME. The MME then asks the base station to allocate radio resources and establish the connection to the UE. During handoff of a UE, the source base station sends the handoff request to the target base station. After receiving an acknowledgement, the source base station transfers the UE state (e.g., buffered packets) to the target base station. The target base station also informs the MME that the UE has changed cells, and the previous base station to release resources (e.g., remove the GTP tunnel).

The S-GW and P-GW are also involved in routing, running protocols such as OSPF. The Policy Control and Charging Function (PCRF) manages flow-based charging in the P-GW. The PCRF also provides the QoS authorization (QoS class identifier and bit rates) that decides how to treat each traffic flow, based on the user's subscription profile. QoS policies can be dynamic, e.g., based on time of day. This must be enforced at the P-GW. The Home Subscriber Server (HSS) contains subscription information for each user, such as the QoS profile, any access restrictions for roaming, and the associated MME. In times of cell congestion, a base station reduces the max rate allowed for subscribers according to their profiles, in coordination with the P-GW.

Today's cellular network architectures have several major limitations. Centralizing data-plane functions such as monitoring, access control, and quality-of-service functionality at the packet gateway introduces scalability challenges. This

problem

978-0-7695-4870-8/12 $26.00 © 2012 IEEE
DOI 10.1109/EWSDN.2012.28

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0033

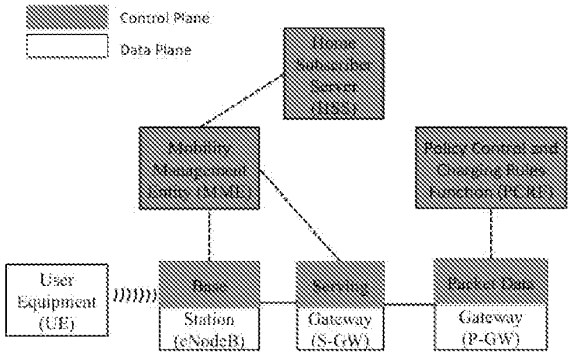


Fig. 2. Simplified LTE network architecture

makes the equipment very expensive (e.g., more than 6 million dollars for a Cisco packet gateway). Centralizing data-plane functions at the cellular-Internet boundary forces all traffic through the P-GW, including traffic between users on the same cellular network, making it difficult to host popular content *inside* the cellular network. In addition, the network equipment has vendor-specific configuration interfaces, and communicate through complex control-plane protocols, with a large and growing number of tunable parameters (e.g., several thousand parameters for base stations). As such, carriers have (at best) indirect control over the operation of their networks, with little ability to create innovative services.

B. SDN for Cellular Data Networks

Cellular data networks are ripe for the introduction of Software Defined Networking (SDN), where the network equipment performs basic packet-processing functions at the behest of applications running on a logically-centralized controller. In the next section, we discuss how SDN could give cellular operators greater control over their equipment, simplify network management, and introduce value-added services. SDN can enable carriers to distribute data-plane rules over multiple, cheaper network switches, reducing the scalability pressure on the packet gateway and enabling flexible handling of traffic that stays within the cellular network.

Supporting real-time updates to many fine-grained packet-handling rules raises significant scalability challenges. Frequent user mobility can require forwarding state at the level of individual subscribers, and the state must change quickly to avoid service disruptions. To detect when subscribers exceed their usage caps, the switches must perform fine-grain monitoring of traffic volumes. The network must adapt quickly to the measurement data to adapt QoS policies, or transcode content to offer good service during times of congestion. To address these challenges, we propose four main extensions to controllers, switches, and base stations:

Flexible policies on subscriber attributes: Many controller applications need to apply policy based on the properties of cellular subscribers, including the network provider, device type, subscriber type, and recent usage. The controller should automatically translate policies based on subscriber attributes to packet-processing rules based on IP addresses and network locations.

Scalability through local switch agents: The need for fast and frequent updates to a large amount of data-plane state would put tremendous pressure on a central controller. Instead, switches should run software agents that perform simple local actions (such as polling traffic counters and comparing against thresholds), at the behest of the controller.

Flexible switch patterns and actions: Today's OpenFlow already supports flexible packet classification and actions. That said, cellular networks would benefit from support for deep-packet inspection, header compression, and message-based control protocols like SCTP [2].

Remote control of virtualized base-station resources: Virtualizing the base station by time slot and subcarriers can give different applications, traffic classes, or virtual operators the illusion of a dedicated base station. An open API between the controller and base station can enable remote control of radio resource allocation, admission control, handoff, and paging.

For prototyping our architecture, we are encouraged by the availability of the open source LTE base station code provided by openairinterface.org and the open source LTE packet core code provided by openepc.net.

II. Controller Applications

In this section, we identify several major challenges in cellular data networks, and how Software Defined Networking can enable better solutions.

A. Directing Traffic Through Middleboxes

Cellular network operators offer many fine-grained services implemented in network appliances, or middleboxes. In a dynamic environment, cellular providers need to adapt video quality in real time based on cell tower congestion, device type, and the subscriber's service plan [3]. To improve security for enterprise customers, providers direct traffic through intrusion detection and prevention systems. Certain legacy devices need in-network support for echo cancellation for VoIP calls. However, today's cellular carriers do not have fine-grained control over routing, forcing them to either direct excess traffic through unnecessary middleboxes or manage an unwieldy set of tunnels.

SDN provides fine-grained packet classifier and flexible routing, which can easily direct a chosen subset of traffic through a set of middleboxes. As a result, middleboxes will handle much less traffic, making them much cheaper. In addition, with support for deep-packet inspection, SDN switches could support some middlebox functionality directly, reducing the number of extra devices in the network.

B. Monitoring for Network Control & Billing

Due to frequent user mobility, and changing channel conditions, cellular networks are subject to rapid changes in traffic demands, and frequent signaling messages to migrate connections. Real-time traffic monitoring is crucial for triggering fast

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0034

adaptation. This includes load balancing data traffic from base stations to a different S-GW, and from a different S-GW to a different P-GW, load balancing control traffic from the base station and S-GW to a different MME. Real-time monitoring also enables rapid per-application content adaptation (e.g., video conferencing, or streaming from Netflix) to meet per-subscriber QoS. Existing traffic-monitoring solutions [4] require additional equipment that captures every packet at every interface of a S-GW, and provides a summary to a backend SQL server every few minutes. These measurements provide no real-time visibility into the eNodeB, S-GW, and MME.

notify

The packet-handling rules in SDN switches include byte and packet counters. By adjusting these rules over time, the cellular provider can efficiently monitor traffic at different levels of granularity to drive real-time control loops on the SDN controller. In addition, associating packet classifiers with traffic counters is useful to drive billing decisions and determine whether a subscriber has reached a usage cap. The recent interest in allowing content providers to cover usage charges for mobile users will put even more pressure on cellular providers to collect fine-grained measurements.

C. Seamless Subscriber Mobility

Cellular networks must respond quickly to subscriber mobility to avoid disruptions in service. Yet, today's cellular providers do not have direct control over routing, or common protocols for controlling forwarding across different cellular technologies (e.g., 3G, LTE, WiMax, and WiFi). As a result, handoff across technologies involves complex procedures that lead to longer delays and higher packet loss rates.

comm protocols

SDN would provide a common control protocol (e.g., OpenFlow) that works across different cellular technologies, making mobility management much easier. In addition, rather than performing hop-by-hop signaling to create a new session, the controller can push new forwarding rules to multiple switches at the same time for a lower set-up delay.

D. QoS and Access Control Policies

In today's networks, the packet gateway is the central point for fine-grained policy enforcement and charging based on the subscriber profile, application, and usage. The P-GW classifies packets based on the 5-tuple of the TCP/IP header and either drops packets (if they violate a firewall policy) or map them into QoS classes. These QoS classes further map into the Diff-Serve Code Points (DSCP) when the packets traverses IP networks en route to the base station. At the base station, only simple policies such as a maximum rate are enforced. This leads to scalability problems at the P-GW, and missed opportunities to optimize the use of bandwidth inside the cellular network.

SDN would enable the distributed enforcement of QoS and firewall policies based on a network-wide view. Distributed enforcement is especially important for handling any traffic that stays within the cellular network. A controller application running on the controller can spread access-control rules over multiple switches, and manage the scheduling of traffic by QoS classes across multiple hops in the network.

E. Virtual Cellular Operators

Today's cellular networks have relatively limited support for virtualization. LTE can isolate different enterprise customers' traffic into virtual private networks using traditional BGP/MPLS VPN technologies. However, LTE does not allow different carriers to share the infrastructure to offer a complete virtual LTE network to their customers. Virtual operators may want to innovate in mobility management, policy, and charging, without investing the substantial resources necessary to build and manage a wireless network. For example, content providers like Akamai could leverage a virtual infrastructure to better deliver content to mobile users.

Virtualization would also be useful to provide isolation and separate control for different classes of traffic. For example, a carrier may want to carry traffic for roaming subscribers on a different virtual network from its own customers, for security reasons.

SDN makes it relatively easy to support network virtualization by partitioning the "flow space" of packet headers. Different controller applications can manage rules acting on each portion of flow space, enabling customized control while ensuring isolation [5]. Virtualizing the cellular network requires virtualizing the base stations [6], [7], by slicing resources at the physical layer (physical channels), link layer (scheduling), or network layer (traffic shaping).

F. Inter-Cell Interference Management

In LTE, every base station can use all subcarriers. However, base stations need to coordinate their subcarrier allocations to avoid harmful interference among users. Currently, inter-cell interference management is done in a distributed fashion. A base station tells its neighboring base stations if it needs to use a higher power to transmit to a UE using a subset of subcarriers. Alternatively, upon perceiving high interference on certain subcarriers, the base station can notify its neighboring base stations to lower their transmission power. Inter-cell interference management algorithm is related to graph coloring algorithm. Due to the lack of a global view, distributed graph coloring algorithms with a limited number of rounds are highly suboptimal [8].

SDN enables centralized control of base stations. A SDN controller will have a global view of the current power and subcarrier allocation profile of base stations. In addition, a SDN controller running on a commodity server would have much more computing resources than most base stations. As a result, a SDN controller can make a more efficient allocation of radio resources to handle new users.

III. Cellular SDN Architecture

Cellular networks need an SDN architecture that offers fine-grain, real-time control without sacrificing scalability. Working our way down from the controller platforms to the base stations, we propose four main extensions to enable SDN in cellular networks. The cellular SDN architecture is shown in Figure 3. First, controller applications should be able to express policy in terms of subscriber attributes, rather than IP addresses or physical locations, as captured in a subscriber

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0035

information base. Second, to improve control-plane scalability, each switch should run a local control agent that performs simple actions (such as polling traffic counters and comparing against a threshold), at the behest of the controller. Third, switches should support more flexible data-plane functionality, such as deep packet inspection and header compression. Fourth, base stations should support remote control of virtualized wireless resources to enable flexible cell management.

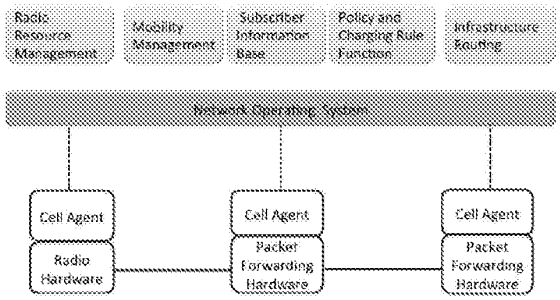


Fig. 3. Cellular SDN architecture

A. Controller: Policies on Subscriber Attributes

The SDN controller consists of a Network Operating System (NOS) running a collection of application modules, such as radio resource management, mobility management, and routing. The handling of individual packets often depends on multiple modules. For example, the flow of traffic through the network depends on the subscriber's location (determined by the mobility manager) and the paths between pairs of network elements (determined by infrastructure routing), and traffic monitoring and packet scheduling depend on the policy and charging rule function. As such, the NOS should support *composition* to combine the results of multiple modules into a single set of packet-handling rules in each switch [9].

Many of the controller application modules need to apply policy based on the properties of cellular subscribers, including the cellular network provider (e.g., whether the user is roaming or not), device type (e.g., whether the user has a legacy phone that requires echo cancellation), subscriber type (e.g., usage cap, parental controls, etc.), and recent usage (e.g., whether the user is near exceeding the usage cap). Yet, the switches match packets and perform actions based on packet header fields, based on ephemeral identifiers such as a subscriber's current IP address and location. To bridge the gap, the controller can maintain a Subscriber Information Base (SIB) that stores and maintains subscriber information, including relatively static subscriber attributes as well as dynamic data like the user's current IP address, location, and total traffic consumption.

The NOS can translate policies expressed in terms of subscriber attributes into switch rules that match on packet headers. Similarly, the NOS can translate network measurements (such as traffic counters) to the appropriate (sets of) subscribers, to allow the application modules to focus on subscribers and their attributes rather than ephemeral network identifiers[1]. The controller can also dynamically divide the network into "slices" that handle all traffic matching some predicate on the subscriber attributes. This allows the cellular provider to isolate roaming traffic, isolate dumb phone traffic or phone traffic using legacy protocols. To enable scalable slicing of semantic space, the controller can instruct ingress switches to mark incoming packets (e.g., using an MPLS label or VLAN tag) sent to or from subscribers with particular attributes.

B. Switch Software: Local Control Agents

Cellular data networks face significant scalability challenges, in terms of the number of subscribers, frequent changes in user location, fine-grain access-control and quality-of-service policies, and real-time adaptation to network conditions. For example, the switches may need to direct a video stream through a transcoding proxy if the network becomes congested, or give certain traffic lower priority if the user exceeds his usage cap. These measurement and control functions could easily overwhelm a logically-centralized controller. In addition, the controller may not be able to respond as quickly to local events as the underlying switches themselves. This argues for having some control-plane functionality on the underlying switches, though arguably not the same complex software that runs on the switches in today's cellular networks.

In particular, each switch can run an agent that performs simple local actions, under the command of the controller. For example, the controller could offload simple measurement tasks to the local agents, such as periodically polling the traffic counters and notifying the controller if a counter exceeds a threshold. The local agent could also perform simple control operations, such as automatically changing the weight or priority of a queue when traffic counts exceed a threshold, or pushing a tag on a packet to direct the traffic through an intermediate middlebox. Performing these operations on the local agents would reduce the load on the controller, and enable faster responses to critical events.

event & notify

Supporting local agents on the switches raises many interesting research problems. Partitioning functionality between the controller and the agents requires ways for application modules to expose the inherent parallelism across agents and the necessary aggregation of information at the controller. In addition, the partitioning of functionality must work in the presence of multiple modules that form a single application. Network updates must be kept consistent amid the partitioning of functionality and user equipment mobility. We plan to explore the design of the local agent and techniques for partitioning functionality in our future work.

C. Switch Hardware: Flexible Packet Processing

Today's OpenFlow [11] switches already support many features needed in cellular networks. Flexible packet classification based on Ethernet, IP, and TCP and UDP header fields enables

[1]This is similar to the way Ethane [10] supports access-control policies based on named principals, rather than IP addresses and network locations, though cellular networks must support a much wider range of attributes and control actions.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0036

fine-grain quality-of-service, access control, and monitoring. The forwarding actions in today's OpenFlow switches would enable carriers to direct selective traffic through middleboxes, change the paths to and from a mobile user, and mark and schedule traffic according to QoS policies. The byte and packet counters associated with each rule would support traffic measurement, real-time adaptation based on congestion or exceeding a subscriber's usage cap, and usage-based billing. Still, these switches may need larger rule tables, or more stages of tables, than today's commodity switches to efficiently support fine-grained policies.

That said, software-defined cellular networks would benefit from new switch capabilities. TCP/UDP port numbers are no longer a sufficiently reliable way to identify applications. Instead, support for deep packet inspection (DPI) would enable finer-grain classification based on the application, such as Web, peer-to-peer, video, and VoIP traffic. This is important to divide traffic into separate traffic classes for different packet-scheduling and routing policies, as commonly done in today's cellular networks [12]. DPI would also help support intrusion detection and prevention systems that analyze packet contents to identify malicious traffic.

To support DPI, the OpenFlow protocol must be extended to add and remove rules where the pattern is a regular expression. To avoid having the DPI module inspect every packet, SDN controller can jointly manage the flow table and the DPI rule table. For example, the switch could first match a packet with a flow-table rule, which includes a flag indicating whether the packet should proceed directly to the output port(s) or go through the DPI table. To contain cost, we do not envision that *every* switch would support DPI. Instead, the controller must place packet-classification rules in the appropriate locations in the underlying network.

In addition to DPI, switches could also support techniques like header compression and decompression to reduce the overhead for applications with small packet payloads. For example, VoIP packets are typically small, making the headers a relatively high fraction of the traffic. Compressing these packets before transmission on low-bandwidth links substantially lowers the overhead. VoIP packet sizes range from 20 to 150 bytes, and the combined overhead of the RTP, UDP, and IP headers is 40 bytes. Robust header compression (ROHC) reduces the 40-byte overhead to 1 byte. As with DPI, this functionality may not be available on every switch, but instead mainly on switches with links to and from low-bandwidth regions of the network.

D. Base Station: Remote Control and Virtualization

Remote control: In LTE, base stations participate in distributed control protocols to manage radio resource allocation, session setup/reconfiguration/teardown, handoff, and paging. Radio resources are inherently shared among base stations. The lack of central control makes it difficult to optimize tasks related to radio access. We decouple the control plane from the radio hardware. We envision that that the SDN radio hardware exposes a well-defined API which can be controlled by the control plane. Running Radio Resource Management Module (RRM) on top of a logically-centralized controller makes it much easier to innovate in admission control, radio resource allocation, and interference management. For example, the RRM can redirect a UE to a nearby lightly loaded base station or increase the transmission power of a congested base station. If the base station has multiple antennas, the control plane can decide whether the antennas should be used for boosting signals (combining diversity to boost signals for delay-sensitive applications) or for spatial multiplexing (multiple parallel transmissions). Although 3G base stations are controlled by a central entity, the radio network controller (RNC) couples control-plane and data-plane functionality, including fine-grained tasks like packet scheduling. In contrast, the RRM module should only perform control-plane functions, and instruct the base station to perform any data-plane operations.

Virtualization: Today's SDN virtualization solutions like FlowVisor [5] can share a single switch data plane among multiple virtual networks. Much as a hypervisor resides between software and hardware on a PC, FlowVisor uses OpenFlow as a hardware abstraction layer to sit logically between control and forwarding paths on a network device. FlowVisor is transparent to both the network hardware and the controller managing the virtual networks. FlowVisor defines a slice as a set of flows running on a topology of switches. The virtualization layer enforces strong isolation between slices. Resources that can be sliced are bandwidth, topology, traffic, device CPU, and forwarding tables.

Technologies like FlowVisor can be extended to slice base station resources, e.g., to create virtual base stations. A virtual base station's resource can be a combination of time slots, subcarriers, and power. For example, a virtual base station can be allocated a subset of time slots, and transmission in each time slot can use maximum transmission power allowed or a subset of subcarriers across all time slots and transmission in each time slot uses a fraction of the maximum transmission power. A slice can ask for base stations with a specific MAC protocol. To support base station virtualization, without modifying the physical-layer protocols, the controller can convey high-level information like the identity of the virtual provider through the control plane. This would allow the UE software to display the virtual provider (rather than the physically broadcasted provider information) without requiring explicit broadcasting of the provider information and cell ID.

IV. Related Work

Our work is heavily inspired and follows the high-level vision of OpenRoads (OpenFlow wireless) [7], which is a platform for innovation and realistic deployment of services for wireless networks. OpenRoads is the first software-defined wireless network. It is mainly based on WiFi and offers no special support for cellular networks. In contrast, our work addresses specific cellular network requirements such as real-time session management which runs on top of SCTP instead of TCP, paging, UE (User Equipment) state tracking, policy enforcement, charging and radio resource management.

Our work focuses on the design of a programmable control plane for cellular networks. It can benefit from a pro-

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0037

grammable data plane [13]. A programmable data plane enables physical layer customization of virtual base stations that traditional MAC scheduling based virtualization techniques [6] can not. For example, each LTE virtual base station can pick a different FFT size to balance multipath effects against peak-to-average-power ratio (PAPR) inefficiencies.

V. Conclusion and Future Work

In this paper, we argue that software defined networking can make cellular networks much simpler and easier to manage, introduce new services, and inter-operate with other wireless network technologies and other operator networks. We sketch out changes and extensions to controller platforms, network switches, and base stations to enable software defined cellular networks. The complete design of an SDN architecture for cellular networks is an exciting avenue for future work.

References

[1] S. Sesia, M. Baker, and I. Toufik, *LTE - the UMTS Long Term Evolution: From Theory to Practice*. John Wiley & Sons, 2011.

[2] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson, "Stream control transmission protocol," Oct 2000. RFC 2960.

[3] "Skyfire launches its Rocket 2.0 platform to help mobile operators deal with a growing data tsunami." http://skyfire.com/images/press_releases/skyfirerocketoptimizerpress.pdf.

[4] "Alcatel-Lucent 9900 wireless network guardian." http://www.alcatel-lucent.com/wps/portal/products/detail?LMSG_CABINET=Solution_Product_Catalog&LMSG_CONTENT_FILE=Products/Product_Detail_000590.xml#tabAnchor1.

[5] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "Can the production network be the testbed?," in *OSDI*, Oct 2010.

[6] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, "NVS: A virtualization substrate for WiMAX networks," in *MOBICOM*, pp. 233–244, ACM, 2010.

[7] K.-K. Yap, R. Sherwood, M. Kobayashi, T.-Y. Huang, M. Chan, N. Handigol, N. McKeown, and G. Parulkar, "Blueprint for introducing innovation into wireless mobile networks," in *Workshop on Virtualized Infrastructure Systems and Architectures*, pp. 25–32, 2010.

[8] F. Kuhn and R. Wattenhofer, "On the complexity of distributed graph coloring," in *PODC*, pp. 7–15, 2006.

[9] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story, and D. Walker, "Frenetic: A network programming language," in *ICFP*, Sep 2011.

[10] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. Gude, N. McKeown, and S. Shenker, "Rethinking enterprise network control," *Trans. on Networking*, vol. 17, Aug 2009.

[11] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks," *SIGCOMM CCR*, vol. 38, no. 2, pp. 69–74, 2008.

[12] "Mobile application assurance on the Alcatel-Lucent 7750 service router mobile gateway: Optimize network resources, enrich and personalize user experiences, and monetize the services," 2011. Application Note.

[13] M. Bansal, J. Mehlman, S. Katti, and P. Levis, "OpenRadio: A programmable wireless dataplane," in *Hot Topics in Software Defined Networks*, pp. 109–114, 2012.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0038

Network Functions Virtualisation

An Introduction, Benefits, Enablers, Challenges & Call for Action

OBJECTIVES

This is a non-proprietary white paper authored by network operators.

The key objective for this white paper is to outline the benefits, enablers and challenges for Network Functions Virtualisation (as distinct from Cloud/SDN) and the rationale for encouraging an international collaboration to accelerate development and deployment of interoperable solutions based on high volume industry standard servers.

CONTRIBUTING ORGANISATIONS & AUTHORS

AT&T:	Margaret Chiosi.
BT:	Don Clarke, Peter Willis, Andy Reid.
CenturyLink:	James Feger, Michael Bugenhagen, Waqar Khan, Michael Fargano.
China Mobile:	Dr. Chunfeng Cui, Dr. Hui Deng.
Colt:	Javier Benitez.
Deutsche Telekom:	Uwe Michel, Herbert Damker.
KDDI:	Kenichi Ogaki, Tetsuro Matsuzaki.
NTT:	Masaki Fukui, Katsuhiro Shimano.
Orange:	Dominique Delisle, Quentin Loudier, Christos Kolias.
Telecom Italia:	Ivano Guardini, Elena Demaria, Roberto Minerva, Antonio Manzalini.
Telefonica:	Diego López, Francisco Javier Ramón Salguero.
Telstra:	Frank Ruhl.
Verizon:	Prodip Sen.

PUBLICATION DATE

October 22-24, 2012 at the "SDN and OpenFlow World Congress", Darmstadt-Germany.

This white paper is available at the following link: http://portal.etsi.org/NFV/NFV_White_Paper.pdf

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0039

Contents

Executive Summary 3
Introduction 5
Definition 5
Relationship with Software Defined Networks (SDN) 5
Fields of Application and Use Cases 6
Benefits of Network Functions Virtualisation 8
The Changing Telecoms Industry Landscape 9
Enablers for Network Functions Virtualisation 9
Cloud Computing 9
Industry Standard High Volume Servers 10
Challenges for Network Functions Virtualisation 10
Recommendations/Call for Action 12
Contact Information 13
References 14
Glossary 14

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0040

Executive Summary

Network Operators' networks are populated with a large and increasing variety of proprietary hardware appliances. To launch a new network service often requires yet another variety and finding the space and power to accommodate these boxes is becoming increasingly difficult; compounded by the increasing costs of energy, capital investment challenges and the rarity of skills necessary to design, integrate and operate increasingly complex hardware-based appliances. Moreover, hardware-based appliances rapidly reach end of life, requiring much of the procure-design-integrate-deploy cycle to be repeated with little or no revenue benefit. Worse, hardware lifecycles are becoming shorter as technology and services innovation accelerates, inhibiting the roll out of new revenue earning network services and constraining innovation in an increasingly network-centric connected world.

Network Functions Virtualisation aims to address these problems by leveraging standard IT virtualisation technology to consolidate many network equipment types onto industry standard high volume servers, switches and storage, which could be located in Datacentres, Network Nodes and in the end user premises. We believe Network Functions Virtualisation is applicable to any data plane packet processing and control plane function in fixed and mobile network infrastructures.

We would like to emphasise that we see Network Functions Virtualisation as highly complementary to Software Defined Networking (SDN). These topics are mutually beneficial but are not dependent on each other. Network Functions can be virtualised and deployed without an SDN being required and vice-versa.

Virtualising Network Functions could potentially offer many benefits including, but not limited to:

- Reduced equipment costs and reduced power consumption through consolidating equipment and exploiting the economies of scale of the IT industry.
- Increased speed of Time to Market by minimising the typical network operator cycle of innovation. Economies of scale required to cover investments in hardware-based functionalities are no longer applicable for software-based development, making feasible other modes of feature evolution. Network Functions Virtualisation should enable network operators to significantly reduce the maturation cycle.
- Availability of network appliance multi-version and multi-tenancy, which allows use of a single platform for different applications, users and tenants. This allows network operators to share resources across services and across different customer bases.
- Targeted service introduction based on geography or customer sets is possible. Services can be rapidly scaled up/down as required.
- Enables a wide variety of eco-systems and encourages openness. It opens the virtual appliance market to pure software entrants, small players and academia, encouraging more innovation to bring new services and new revenue streams quickly at much lower risk.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0041

To leverage these benefits, there are a number of technical challenges which need to be addressed:

- Achieving high performance virtualised network appliances which are portable between different hardware vendors, and with different hypervisors.

hardware independent

- Achieving co-existence with bespoke hardware based network platforms whilst enabling an efficient migration path to fully virtualised network platforms which re-use network operator OSS/BSS. OSS/BSS development needs to move to a model in-line with Network Functions Virtualisation and this is where SDN can play a role.
- Managing and orchestrating many virtual network appliances (particularly alongside legacy management systems) while ensuring security from attack and misconfiguration.
- Network Functions Virtualisation will only scale if all of the functions can be automated.
- Ensuring the appropriate level of resilience to hardware and software failures.
- Integrating multiple virtual appliances from different vendors. Network operators need to be able to "mix & match" hardware from different vendors, hypervisors from different vendors and virtual appliances from different vendors without incurring significant integration costs and avoiding lock-in.

The authors of this white paper believe (and have in some cases demonstrated) that solutions to these technical challenges are available, or could be made available, and recommend that the IT and Networks industries combine their complementary expertise and resources in a joint collaborative effort to reach broad agreement on standardised approaches and common architectures which address these technical challenges, and which are interoperable and have economies of scale.

To accelerate progress, a new network operator-led Industry Specification Group (ISG) with open membership is being setup under the auspices of ETSI to work through the technical challenges for Network Functions Virtualisation as outlined in this white paper. The formal creation process of this ETSI ISG has been started and is expected to be completed by mid-November 2012.

In order to chart the way forward for Network Functions Virtualisation the wider industry is asked to provide feedback as detailed in this white paper.

End - Executive Summary.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0042

Introduction

Network operators' networks are populated with a large and increasing variety of proprietary hardware appliances. To launch a new network service often requires yet another variety and finding the space and power to accommodate these boxes is becoming increasingly difficult; compounded by the increasing costs of energy, capital investment challenges and the rarity of skills necessary to design, integrate and operate increasingly complex hardware-based appliances. Moreover, hardware-based appliances rapidly reach end of life, requiring much of the procure-design-integrate-deploy cycle to be repeated with little or no revenue benefit. Worse, hardware lifecycles are becoming shorter as technology and services innovation accelerates, inhibiting the roll out of new revenue earning network services and constraining innovation in an increasingly network-centric connected world.

Definition

Network Functions Virtualisation aims to transform the way that network operators architect networks by evolving standard IT virtualisation technology to consolidate many network equipment types onto industry standard high volume servers, switches and storage, which could be located in Datacentres, Network Nodes and in the end user premises, as illustrated in Figure 1. It involves the implementation of network functions in software that can run on a range of industry standard server hardware, and that can be moved to, or instantiated in, various locations in the network as required, without the need for installation of new equipment.

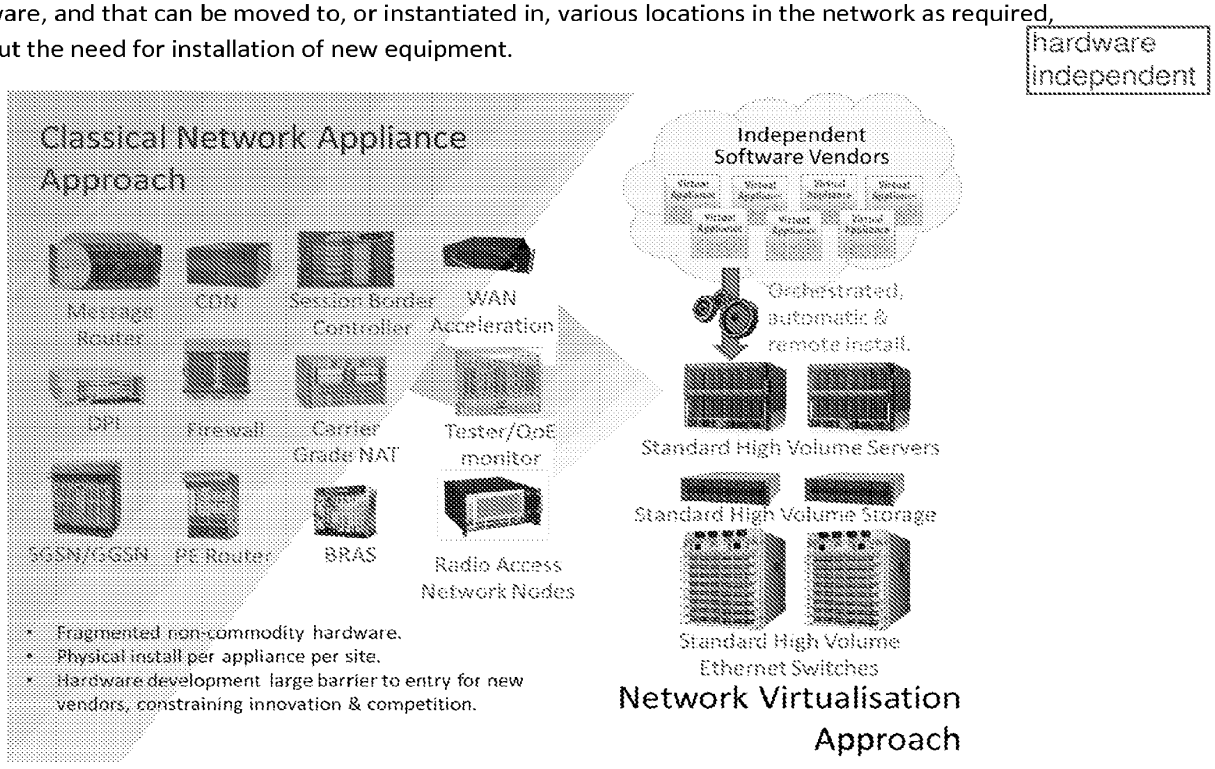


Figure 1: Vision for Network Functions Virtualisation

Relationship with Software Defined Networks (SDN)

As shown in Figure 2, Network Functions Virtualisation is highly complementary to Software Defined Networking (SDN), but not dependent on it (or vice-versa). Network Functions Virtualisation can be

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0043

implemented without a SDN being required, although the two concepts and solutions can be combined and potentially greater value accrued.

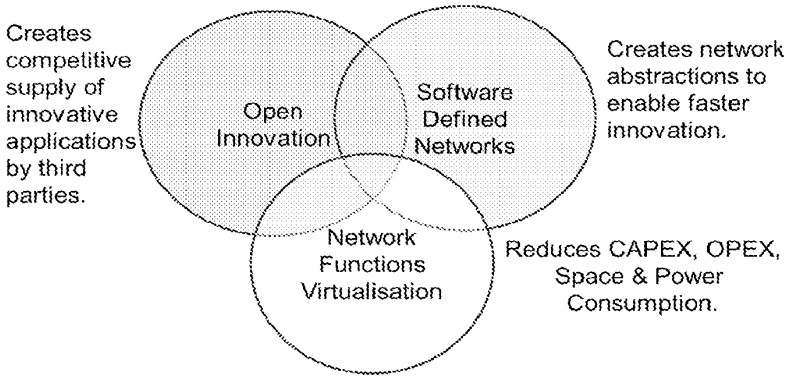


Figure 2: Network Functions Virtualisation Relationship with SDN

Network Functions Virtualisation goals can be achieved using non-SDN mechanisms, relying on the techniques currently in use in many datacentres. But approaches relying on the separation of the control and data forwarding planes as proposed by SDN can enhance performance, simplify compatibility with existing deployments, and facilitate operation and maintenance procedures.

Network Functions Virtualisation is able to support SDN by providing the infrastructure upon which the SDN software can be run. Furthermore, Network Functions Virtualisation aligns closely with the SDN objectives to use commodity servers and switches.

We intend to work closely with organisations progressing work on SDN such as the Open Networking Foundation (ONF) whose work we will specifically take into account.

Fields of Application and Use Cases

Network Functions Virtualisation is applicable to any data plane packet processing and control plane function in mobile and fixed networks. Potential examples that can be listed include (not in any particular order):

- Switching elements: BNG, CG-NAT, routers.
- Mobile network nodes: HLR/HSS, MME, SGSN, GGSN/PDN-GW, RNC, Node B, eNode B.
- Functions contained in home routers and set top boxes to create virtualised home environments.
- Tunnelling gateway elements: IPSec/SSL VPN gateways.
- Traffic analysis: DPI, QoE measurement.
- Service Assurance, SLA monitoring, Test and Diagnostics.
- NGN signalling: SBCs, IMS.
- Converged and network-wide functions: AAA servers, policy control and charging platforms.
- Application-level optimisation: CDNs, Cache Servers, Load Balancers, Application Accelerators.
- Security functions: Firewalls, virus scanners, intrusion detection systems, spam protection.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0044

Further studies are required to identify those network functions for which virtualisation yields most benefits. Use cases of interest to the authors include (not in any particular order):

- A software-based DPI, providing advanced traffic analysis and multi-dimensional reporting, and showing the possibility of making off-the-shelf hardware work at actual line rates. Software-based DPI can be pervasively deployed in the network, providing much better analysis capabilities, as well as simpler mechanisms for deployment, update, testing, and to scale it to changing workloads.
- IP node implementations, supporting - for example, but not limited to: CG-NAT and BRAS capabilities on standard high-end servers, offering the opportunity for an effective re-use of hardware as the demand for such capabilities evolves.
- The virtualisation of services and capabilities that presently require dedicated hardware appliances on customer premises (home environment to small branch office to large corporate premises), including but not restricted to: firewall, web security, IPS/IDS, WAN acceleration and optimisation, and router functions. The virtualisation of the home environment including routers, hubs and set top boxes would potentially enable a simpler and seamless migration to IPv6, reduce energy consumption and avoid successive hardware updates as broadband applications and services evolve.
- The virtualisation of Content Distribution Networks (CDN), with the initial goal of extending and scaling Content Delivery Services more easily, and also with the objective of maximizing hardware re-use in PoPs by being able to install other Service Delivery Applications (e.g. Web Acceleration) on demand. Virtualisation of CDNs will also allow the hosting of CDN services from potential business partners, like external CDN providers.
- The virtualisation of a mobile core network targeting at a more cost efficient production environment, which allows network operators to cope with the increasing traffic demand in mobile networks, and leading to better resource utilization (including energy savings), more flexible network management (no need to change hardware for nodes' upgrades), hardware consolidation, easier multi-tenancy support and faster configuration of new services. Network Functions Virtualisation in mobile networks can also be used to create core network instances optimized for specific services, e.g. for Machine-to-Machine communications (M2M).

distributed applications

- Co-ordinated implementation of cloud and networking for enterprises, allowing on-demand services to be offered and providing capital efficiency for enterprise customers and network operators.
- Hybrid fibre-DSL nodes are located deep in the external network in street cabinets, underground and on poles. These nodes must be very low power consumption and very low/zero maintenance to be economic. Virtualisation could be used to reduce hardware complexity at the remote node, saving energy and providing an enhanced degree of future proofing as services evolve. These remote nodes could more economically provide both fixed and wireless access if key functions were virtualised on a common platform.
- Network Functions Virtualisation can also be used to provide an efficient production environment which can commonly be used by different applications, users and tenants, thus

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0045

supporting the coexistence of several versions and variants of a network service (including test versions and beta versions).

Benefits of Network Functions Virtualisation

We believe the application of Network Functions Virtualisation brings many benefits to network operators, contributing to a dramatic change in the telecommunications industry landscape. Benefits we foresee include (not in any particular order):

- Reduced equipment costs and reduced power consumption through consolidating equipment and exploiting the economies of scale of the IT industry (~9.5M Servers shipped in 2011 compared with ~1.5M routers forecast for 2012).[1,2]
- Increased velocity of Time to Market by minimising the typical network operator cycle of innovation. Economies of scale required to cover investments in hardware-based functionalities are no longer applicable for software-based development, making feasible other modes of feature evolution. Network Functions Virtualisation should enable network operators to significantly reduce the maturation cycle.
- The possibility of running production, test and reference facilities on the same infrastructure provides much more efficient test and integration, reducing development costs and time to market.
- Targeted service introduction based on geography or customer sets is possible. Services can be rapidly scaled up/down as required. In addition, service velocity is improved by provisioning remotely in software without any site visits required to install new hardware.
- Enabling a wide variety of eco-systems and encouraging openness. It opens the virtual appliance market to pure software entrants, small players and academia, encouraging more innovation to bring new services and new revenue streams quickly at much lower risk.
- Optimizing network configuration and/or topology in near real time based on the actual traffic/mobility patterns and service demand. For example, optimisation of the location & assignment of resources to network functions automatically and in near real time could provide protection against failures without engineering full 1+1 resiliency.
- Supporting multi-tenancy thereby allowing network operators to provide tailored services and connectivity for multiple users, applications or internal systems or other network operators, all co-existing on the same hardware with appropriate secure separation of administrative domains.
- Reduced energy consumption by exploiting power management features in standard servers and storage, as well as workload consolidation and location optimisation. For example, relying on virtualisation techniques it would be possible to concentrate the workload on a smaller number of servers during off-peak hours (e.g. overnight) so that all the other servers can be switched off or put into an energy saving mode.[3]
- Improved operational efficiency by taking advantage of the higher uniformity of the physical network platform and its homogeneity to other support platforms:

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0046

- IT orchestration mechanisms provide automated installation, scaling-up and scaling-out of capacity, and re-use of Virtual Machine (VM) builds.[4]
- Eliminating the need for application-specific hardware. The skills base across the industry for operating standard high volume IT servers is much larger and less fragmented than for today's telecom-specific network equipment.
- Reduction in variety of equipment for planning & provisioning. Assuming tools are developed for automation and to deal with the increased software complexity of virtualisation.
- Option to temporarily repair failures by automated re-configuration and moving network workloads onto spare capacity using IT orchestration mechanisms. This could be used to reduce the cost of 24/7 operations by mitigating failures automatically.
- The potential to gain more efficiency between IT and Network Operations.
- The potential to support in-service software upgrade (ISSU) with easy reversion by installing the new version of a Virtualised Network Appliance (VNA) as a new Virtual Machine (VM). Assuming traffic can be transferred from the old VM to the new VM without interrupting service. For some applications it may be necessary to synchronise the state of the new VM with the old VM.

substantially no interruption to the operation

The Changing Telecoms Industry Landscape

Although Network Functions Virtualisation brings many advantages to the telecommunications industry it is likely to transform the vendor landscape. Each player will need to position/re-position itself in the new Network Functions Virtualisation market.

This is not as disruptive as it may seem because network equipment vendors already implement some of their solutions by combining their proprietary software with industry standard hardware and software components, but in a proprietary way. Enabling their proprietary software to run on industry standard hardware in a standardised way may be a significant opportunity for existing players because their software and networking know-how is where the real value is in many cases. Some major industry players are already moving in this direction by offering virtualised versions of their products. The challenge for network operators is how to migrate their operations and skill base to a software based networking environment while carefully re-targeting investment to maximise re-use of existing systems and processes.

Enablers for Network Functions Virtualisation

Several recent technology developments make the goals of Network Functions Virtualisation achievable. This section describes these enablers and briefly discusses relevance.

Cloud Computing

Network Functions Virtualisation will leverage modern technologies such as those developed for cloud computing. At the core of these cloud technologies are virtualisation mechanisms: hardware virtualisation by means of hypervisors, as well as the usage of virtual Ethernet switches (e.g. vswitch) for connecting traffic between virtual machines and physical interfaces. For communication-oriented

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0047

functions, high-performance packet processing is available through high-speed multi-core CPUs with high I/O bandwidth, the use of smart Ethernet NICs for load sharing and TCP Offloading, and routing packets directly to Virtual Machine memory, and poll-mode Ethernet drivers (rather than interrupt driven, for example Linux NAPI and Intel's DPDK).

Cloud infrastructures provide methods to enhance resource availability and usage by means of orchestration and management mechanisms, applicable to the automatic instantiation of virtual appliances in the network, to the management of resources by assigning virtual appliances to the correct CPU core, memory and interfaces, to the re-initialisation of failed VMs, to snapshot VM states and the migration of VMs.

programming using upper layer APIs

Finally, the availability of open APIs for management and data plane control, like OpenFlow, OpenStack, OpenNaaS or OGF's NSI, provide an additional degree of integration of Network Functions Virtualisation and cloud infrastructure.

Industry Standard High Volume Servers

The use of industry standard high volume servers is a key element in the economic case for Network Functions Virtualisation. Network Functions Virtualisation leverages the economies of scale of the IT industry. An industry standard high volume server is a server built using standardised IT components (for example x86 architecture) and sold in the millions. A common feature of industry standard high volume servers is that there is competitive supply of the subcomponents which are interchangeable inside the server.

We believe that Network Appliances which depend on the development of bespoke Application Specific Integrated Circuits (ASICs) will become increasingly uncompetitive against general purpose processors as the cost of developing ASICs increases exponentially with decreasing feature size.[5] Merchant silicon will still be applicable for commodity functions implemented at scale, and ASICs will still be applicable for some types of very high throughput applications.

Challenges for Network Functions Virtualisation

There are a number of challenges to implement Network Functions Virtualisation which need to be addressed by the community interested in accelerating progress. How this effort could be progressed is described later in this document. Challenges we have identified are (not in any particular order):-

- **Portability/Interoperability.** The ability to load and execute virtual appliances in different but standardised datacentre environments, provided by different vendors for different operators. The challenge is to define a unified interface which clearly decouples the software instances from the underlying hardware, as represented by virtual machines and their hypervisors. Portability and Interoperability is very important as it creates different ecosystems for virtual appliance vendors and datacentre vendors, while both ecosystems are clearly coupled and depend on each other. Portability also allows the operator the freedom to optimise the location and required resources of the virtual appliances without constraints.
- **Performance Trade-Off.** Since the Network Functions Virtualisation approach is based on industry standard hardware (i.e. avoiding any proprietary hardware such as acceleration

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0048

engines) a probable decrease in performance has to be taken into account. The challenge is how to keep the performance degradation as small as possible by using appropriate hypervisors and modern software technologies, so that the effects on latency, throughput and processing overhead are minimized. The available performance of the underlying platform needs to be clearly indicated, so that virtual appliances know what they can get from the hardware. The authors of the white paper believe that using the right technology choice will allow virtualisation not only of network control functions but also data/user plane functions.

- **Migration and co-existence of legacy & compatibility with existing platforms**. Implementations of Network Functions Virtualisation must co-exist with network operators' legacy network equipment and be compatible with their existing Element Management Systems, Network Management Systems, OSS and BSS, and potentially existing IT orchestration systems if Network Functions Virtualisation orchestration and IT orchestration are to converge. The Network Functions Virtualisation architecture must support a migration path from today's proprietary physical network appliance based solutions to more open standards based virtual network appliance solutions. In other words, Network Functions Virtualisation must work in a hybrid network composed of classical physical network appliances and virtual network appliances. Virtual appliances must therefore use existing North Bound Interfaces (for management & control) and interwork with physical appliances implementing the same functions.
- **Management and Orchestration.** A consistent management and orchestration architecture is required. Network Functions Virtualisation presents an opportunity, through the flexibility afforded by software network appliances operating in an open and standardised infrastructure, to rapidly align management and orchestration North Bound Interfaces to well defined standards and abstract specifications. This will greatly reduce the cost and time to integrate new virtual appliances into a network operator's operating environment. Software Defined Networking (SDN) further extends this to streamlining the integration of packet and optical switches into the system e.g. a virtual appliance or Network Functions Virtualisation orchestration system may control the forwarding behaviours of physical switches using SDN.
- **Automation.** Network Functions Virtualisation will only scale if all of the functions can be automated. Automation of process is paramount to success.
- **Security & Resilience.** Network operators need to be assured that the security, resilience and availability of their networks are not impaired when virtualised network functions are introduced. Our initial expectations are that Network Functions Virtualisation improves network resilience and availability by allowing network functions to be recreated on demand after a failure. A virtual appliance should be as secure as a physical appliance if the infrastructure, especially the hypervisor and its configuration, is secure. Network operators will be seeking tools to control and verify hypervisor configurations. They will also require security certified hypervisors and virtual appliances.
- **Network Stability.** Ensuring stability of the network is not impacted when managing and orchestrating a large number of virtual appliances between different hardware vendors and hypervisors. This is particularly important when, for example, virtual functions are relocated, or during re-configuration events (e.g. due to hardware and software failures) or due to cyber-attack. This challenge is not unique to Network Functions Virtualisation. Potential instability might also occur in current networks, depending on unwanted combinations of diverse control and optimization mechanisms, for example acting on either the underlying transport network or on the higher layers' components (e.g. flow admission control, congestion control, dynamic routing and allocations, etc.). It should be noted that

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0049

occurrence of network instability may have primary effects, such as jeopardizing, even dramatically, performance parameters or compromising an optimized use of resources. Mechanisms capable of ensuring network stability will add further benefits to Network Functions Virtualisation.

- **Simplicity.** Ensuring that virtualised network platforms will be simpler to operate than those that exist today. A significant and topical focus for network operators is simplification of the plethora of complex network platforms and support systems which have evolved over decades of network technology evolution, while maintaining continuity to support important revenue generating services. It is important to avoid trading one set of operational headaches for a different but equally intractable set of operational headaches.
- **Integration.** Seamless integration of multiple virtual appliances onto existing industry standard high volume servers and hypervisors is a key challenge for Network Functions Virtualisation. Network operators need to be able to "mix & match" servers from different vendors, hypervisors from different vendors and virtual appliances from different vendors without incurring significant integration costs and avoiding lock-in. The ecosystem must offer integration services and maintenance and third-party support; it must be possible to resolve integration issues between several parties. The ecosystem will require mechanisms to validate new Network Functions Virtualisation products. Tools must be identified and/or created to address these issues.

Recommendations/Call for Action

Network Functions Virtualisation is already occurring. In a few years, we can expect the communications industry to look and feel similar to the IT industry. There will be a wider range of business models more suited to a software industry. Operations complexity will be abstracted away by more automation and self-provisioning will be more common. As detailed in this white paper, Network Functions Virtualisation will deliver many benefits for network operators and their partners and customers whilst offering the opportunity to create new types of eco-systems (alongside traditional supply models based on preferred strategic partners) which will encourage and support rapid innovation with reduced cost and reduced risk. To reap these benefits the technical challenges, as described above, must be addressed by the industry.

The authors of this white paper believe that solutions to these technical challenges are available (or could be made available) and recommend that the IT and Network industries combine their complementary expertise and resources in a joint collaborative effort to reach broad agreement on standardised approaches and common architectures which address these technical challenges, and which are interoperable and have economies of scale.

To accelerate progress, a new network operator-led Industry Specification Group (ISG) with open membership is being setup under the auspices of ETSI to work through the technical challenges for Network Functions Virtualisation as outlined in this white paper. The formal creation process of this ETSI ISG has been started and is expected to be completed by mid-November 2012.

While ETSI will provide the organisation for this initiative, the ISG will downstream its work to all relevant organisations and seek to encourage convergence of IT and Network Standardisation efforts in this space.

The key objectives, activities and outputs of the ISG are expected to be:

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0050

- The ISG will initiate joint studies which address the technical challenges for Network Functions Virtualisation as documented in this white paper. With consensus outputs detailed in white papers openly published to the wider industry. It is expected that study items would be identified in a "Network Operator Requirements" white paper prepared by the ISG as an initial deliverable.
- Although parented under ETSI, the ISG will not be a Standards Development Organisation (SDO) in itself; it will reference the work of existing SDOs where applicable. Where there are gaps, it will encourage member contributions to SDOs requesting them to take into account the recommendations of the ISG (i.e. requesting updates to existing standards, or to create new standards if appropriate).
- The ISG will identify other expert bodies that could assist with this work, and request that they undertake studies related to their scope to help move the industry forward.
- The ISG will facilitate members to share technical learning in a collaborative way which does not compromise confidentiality.
- The first meeting of the ISG will take place in January 2013. Details will be announced as soon as possible after publication of this white paper.

In order to chart the way forward for Network Functions Virtualisation, the wider industry is asked to provide feedback:

- **Network Operators:** To state their interest in joining the ISG, thus subscribing to the targets and approach as outlined in this white paper.
- **Existing organizations and forums:** To state if they see themselves being impacted by the work addressed in the ISG, and to provide information on how they would like to interact with it.
- **Vendors:** To state their interest in joining the ISG. Vendors are furthermore requested to state the study items in which they are most interested.

Please provide your input to the originators of this white paper as soon as possible and no later than 16th November 2012. Contact details below.

Contact Information

Please provide your feedback to any of the following. Please be clear on what you are asking us to do with the information provided. **Your input must not be confidential.**

AT&T:	Margaret Chiosi, mchiosi@att.com
BT:	Don Clarke, don.clarke@bt.com Peter Willis, peter.j.willis@bt.com
CenturyLink:	James Feger, james.feger@centurylink.com Michael Bugenhagen, michael.k.bugenhagen@centurylink.com Waqar Khan, waqar.khan@centurylink.com Michael Fargano, michael.fargano@centurylink.com

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0051

China Mobile:	Clark Chen, chenkuilin@chinamobile.com Jinri Huang, huangjinri@chinamobile.com
Colt:	Javier Benitez, javier.benitez@colt.net
Deutsche Telekom:	Uwe Michel, uwe.michel@telekom.de Herbert Damker, herbert.damker@telekom.de
KDDI:	Kenichi Ogaki, ke-oogaki@kddi.com
NTT:	Masaki Fukui, fukui.masaki@lab.ntt.co.jp Katsuhiro Shimano, shimano.katsuhiro@lab.ntt.co.jp
Orange:	Dominique Delisle, dominique.delisle@orange.com Quentin Loudier, quentin.loudier@orange.com Christos Kolias, christos.kolias@orange.com
Telecom Italia:	Ivano Guardini, ivano.guardini@telecomitalia.it Elena Demaria, elena.demaria@telecomitalia.it
Telefonica:	Diego López, diego@tid.es Francisco Javier Ramón Salguero, fjrs@tid.es
Telstra:	Frank Ruhl, frank.ruhl@team.telstra.com
Verizon:	Prodip Sen, prodip.sen@verizon.com

References

1. Server shipments (9.5M units in 2011): http://www.crn.com/news/data-center/232601815/gartner-2011-server-sales-grow-but-hard-drive-issues-take-a-toll.htm
2. Router shipments (1.5M units forecast for year 2012): http://www.isuppli.com/Manufacturing-and-Pricing/MarketWatch/Pages/EMS-Router-Shipments-Backed-by-Market-Growth-in-2012.aspx
3. http://labs.chinamobile.com/cran/wp-content/uploads/CRAN_white_paper_v2_5_EN.pdf
4. http://en.wikipedia.org/wiki/Orchestration_(computing)
5. http://www.edac.org/downloads/resources/profitability/HandelJonesReport.pdf

Glossary

AAA	Authentication-Authorisation-Accounting
API	Application Programming Interface (a set of routines, protocols, and tools for building software applications)
ASIC	Application Specific Integrated Circuit
BNG	Broadband Network Gateway
BRAS	Broadband Remote Access Server
BSS	Business Support System
CDN	Content Distribution Network
CG-NAT/NAT	Carrier Grade Network Address Translation
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Service
DPDK	Data Plane Development Kit

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0052

DPI	Deep Packet Inspection
DSL	Digital Subscriber Line (typically ADSL2/+, VDSL2, G.fast)
Ecosystem	Equipment and software vendors, integration services vendors, network operators, service providers, academia, etc.
EMS	Element Management System
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute (in practise ETSI has global membership, including the Americas and Asia) .
HSS	Home Subscriber Service
HV	Hypervisor
I/O	Input/Output
IMS	IP Multimedia System
Industry Standard High Volume Server	A server built using standardised IT components (for example x86 architecture) and sold in the millions.
IPR	Intellectual Property Rights
IPSec/SSL	IP Security/ Secure Sockets Layer
IPv6	Internet Protocol version 6
ISG	Industry Specification Group. An organisation under the auspices of the European Telecommunications Standards Institute (ETSI).
ISSU	In Service Software Upgrade
IT	Information Technology
KVM	Kernel based Virtual Machine
LAN	Local Area Network
LB	Load Balancer
M2M	Machine-to-Machine communications
MME	Mobility Management Entity
MVNO	Mobile Virtual Network Operator
NFV	Network Functions Virtualisation
NGN	Next Generation Network
NIC	Network Interface Controller
NMS	Network Management System
OAM	Operations Administration & Maintenance
OGF	Open Grid Forum
ONF	Open Networking Foundation
OpenFlow	Specifications developed by the Open Networking Foundation
OpenNaaS	Specifications developed by the OpenNaaS community
OpenStack	Specifications developed by the OpenStack Foundation
OSS	Operations Support System
PDN-GW	Packet Data Network Gateway
PE Router	Provider Edge Router
PoP	Point of Presence
QoE	Quality of Experience
SBC	Session Border Controller
SDN	Software Defined Network
SDO	Standards Development Organisation
SGSN/GGSN	Serving GPRS support node/Gateway GPRS support node
SLA	Service Level Agreement

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0053

TCP	Transmission Control Protocol
VM	Virtual Machine
VMWare	Proprietary Hypervisor
VNA	Virtualised Network Appliance
VPN	Virtual Private Network
vswitch	Any Ethernet switch implemented in software alongside or inside a hypervisor. There are proprietary and open implementations of vswitch.
WAN	Wide Area Network
Xen	Proprietary Hypervisor

END

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0054

OpenFlow: Enabling Innovation in Campus Networks

March 14, 2008

Nick McKeown
Stanford University

Tom Anderson
University of Washington

Hari Balakrishnan
MIT

Guru Parulkar
Stanford University

Larry Peterson
Princeton University

Jennifer Rexford
Princeton University

Scott Shenker
University of California, Berkeley

Jonathan Turner
Washington University in St. Louis

ABSTRACT

This whitepaper proposes OpenFlow: a way for researchers to run experimental protocols in the networks they use every day. OpenFlow is based on an Ethernet switch, with an internal flow-table, and a standardized interface to add and remove flow entries. Our goal is to encourage networking vendors to add OpenFlow to their switch products for deployment in college campus backbones and wiring closets. We believe that OpenFlow is a pragmatic compromise: on one hand, it allows researchers to run experiments on heterogeneous switches in a uniform way at line-rate and with high port-density; while on the other hand, vendors do not need to expose the internal workings of their switches. In addition to allowing researchers to evaluate their ideas in real-world traffic settings, OpenFlow could serve as a useful campus component in proposed large-scale testbeds like GENI. Two buildings at Stanford University will soon run OpenFlow networks, using commercial Ethernet switches and routers. We will work to encourage deployment at other schools; and We encourage you to consider deploying OpenFlow in your university network too.

Categories and Subject Descriptors

C.2 [**Internetworking**]: Routers

General Terms

Experimentation, Design

Keywords

Ethernet switch, virtualization, flow-based

1. THE NEED FOR PROGRAMMABLE NETWORKS

Networks have become part of the critical infrastructure of our businesses, homes and schools. This success has been both a blessing and a curse for networking researchers; their work is more relevant, but their chance of making an impact is more remote. The reduction in real-world impact of any given network innovation is because the enormous installed base of equipment and protocols, and the reluctance to experiment with production traffic, which have created an exceedingly high barrier to entry for new ideas. Today, there is almost no practical way to experiment with new network protocols (e.g., new routing protocols, or alternatives to IP) in sufficiently realistic settings (e.g., at scale carrying real traffic) to gain the confidence needed for their widespread deployment. The result is that most new ideas from the networking research community go untried and untested; hence the commonly held belief that the network infrastructure has "ossified".

Having recognized the problem, the networking community is hard at work developing programmable networks, such as GENI [1] a proposed nationwide research facility for experimenting with new network architectures and distributed systems. These programmable networks call for programmable switches and routers that (using *virtualization*) can process packets for multiple isolated experimental networks simultaneously. For example, in GENI it is envisaged that a researcher will be allocated a *slice* of resources across the whole network, consisting of a portion of network links, packet processing elements (e.g. routers) and end-hosts; researchers program their slices to behave as they wish. A slice could extend across the backbone, into access networks, into college campuses, industrial research labs, and include wiring closets, wireless networks, and sensor networks.

Virtualized programmable networks could lower the barrier to entry for new ideas, increasing the rate of innovation in the network infrastructure. But the plans for nationwide facilities are ambitious (and costly), and it will take years for them to be deployed.

This whitepaper focuses on a shorter-term question closer to home: *As researchers, how can we run experiments in our campus networks?* If we can figure out how, we can start soon and extend the technique to other campuses to benefit the whole community.

To meet this challenge, several questions need answering, including: In the early days, how will college network administrators get comfortable putting experimental equipment (switches, routers, access points, etc.) into their network? How will researchers control a portion of their local network in a way that does not disrupt others who depend on it? And exactly what functionality is needed in network switches to enable experiments? Our goal here is to propose a new switch feature that can help extend programmability into the wiring closet of college campuses.

One approach - that we do not take - is to persuade

commercial "name-brand" equipment vendors to provide an open, programmable, virtualized platform on their switches and routers so that researchers can deploy new protocols, while network administrators can take comfort that the equipment is well supported. This outcome is very unlikely in the short-term. Commercial switches and routers do not typically provide an open software platform, let alone provide a means to virtualize either their hardware or software. The practice of commercial networking is that the standardized external interfaces are narrow (i.e., just packet forwarding), and all of the switch's internal flexibility is hidden. The internals differ from vendor to vendor, with no standard platform for researchers to experiment with new ideas. Further, network equipment vendors are understandably nervous about opening up interfaces inside their boxes: they have spent years deploying and tuning fragile distributed protocols and algorithms, and they fear that new experiments will bring networks crashing down. And, of course, open platforms lower the barrier-to-entry for new competitors.

A few open software platforms already exist, but do not have the performance or port-density we need. The simplest example is a PC with several network interfaces and an operating system. All well-known operating systems support routing of packets between interfaces, and open-source implementations of routing protocols exist (e.g., as part of the Linux distribution, or from XORP [2]); and in most cases it is possible to modify the operating system to process packets in almost any manner (e.g., using Click [3]). The problem, of course, is performance: A PC can neither support the number of ports needed for a college wiring closet (a fanout of 100+ ports is needed per box), nor the packet-processing performance (wiring closet switches process over 100Gbits/s of data, whereas a typical PC struggles to exceed 1Gbit/s; and the gap between the two is widening).

Existing platforms with specialized hardware for line-rate processing are not quite suitable for college wiring closets either. For example, an ATCA-based virtualized programmable router called the Supercharged PlanetLab Platform [4] is under development at Washington University, and can use network processors to process packets from many interfaces simultaneously at line-rate. This approach is promising in the long-term, but for the time being is targeted at large switching centers and is too expensive for widespread deployment in college wiring closets. At the other extreme is NetFPGA [5] targeted for use in teaching and research labs. NetFPGA is a low-cost PCI card with a user-programmable FPGA for processing packets, and 4-ports of Gigabit Ethernet. NetFPGA is limited to just four network interfaces—insufficient for use in a wiring closet.

Thus, the commercial solutions are too closed and inflexible, and the research solutions either have insufficient performance or fanout, or are too expensive. It seems unlikely that the research solutions, with their complete generality, can overcome their performance or cost limitations. A more promising approach is to compromise on generality and to seek a degree of switch flexibility that is:

- Amenable to high-performance and low-cost implementations.
- Capable of supporting a broad range of research.
- Assured to isolate experimental traffic from production traffic.

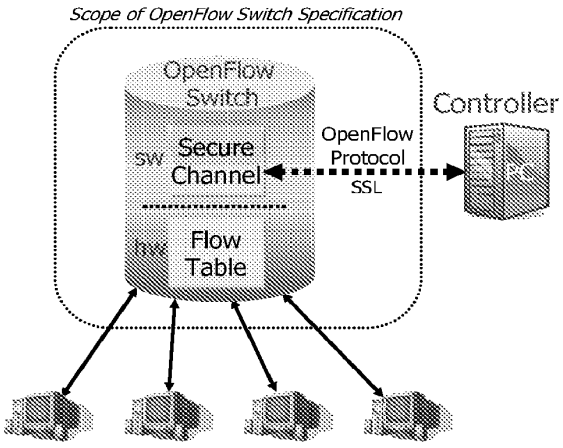


Figure 1: Idealized OpenFlow Switch. The Flow Table is controlled by a remote controller via the Secure Channel.

- Consistent with vendors' need for closed platforms.

This paper describes the OpenFlow Switch a specification that is an initial attempt to meet these four goals.

2. THE OPENFLOW SWITCH

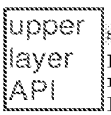
The basic idea is simple: we exploit the fact that most modern Ethernet switches and routers contain flow-tables (typically built from TCAMs) that run at line-rate to implement firewalls, NAT, QoS, and to collect statistics. While each vendor's flow-table is different, we've identified an interesting common set of functions that run in many switches and routers. OpenFlow exploits this common set of functions.

OpenFlow provides an open protocol to program the flow-table in different switches and routers. A network administrator can partition traffic into production and research flows. Researchers can control their own flows - by choosing the routes their packets follow and the processing they receive. In this way, researchers can try new routing protocols, security models, addressing schemes, and even alternatives to IP. On the same network, the production traffic is isolated and processed in the same way as today.

The datapath of an OpenFlow Switch consists of a Flow Table, and an action associated with each flow entry. The set of actions supported by an OpenFlow Switch is extensible, but below we describe a minimum requirement for all switches. For high-performance and low-cost the datapath must have a carefully prescribed degree of flexibility. This means forgoing the ability to specify arbitrary handling of each packet and seeking a more limited, but still useful, range of actions. Therefore, later in the paper, define a basic required set of actions for all OpenFlow switches.

An OpenFlow Switch consists of at least three parts: (1) A *Flow Table*, with an action associated with each flow entry, to tell the switch how to process the flow, (2) A *Secure Channel* that connects the switch to a remote control process (called the *controller*), allowing commands and packets

to be sent between a controller and the switch using (3) The *OpenFlow Protocol*, which provides an open and standard way for a controller to communicate with a switch. By specifying a standard interface (the OpenFlow Protocol) through which entries in the Flow Table can be defined externally, the OpenFlow Switch avoids the need for researchers to program the switch.



It is useful to categorize switches into dedicated OpenFlow switches that do not support normal Layer 2 and Layer 3 processing, and OpenFlow-enabled general purpose commercial Ethernet switches and routers, to which the OpenFlow Protocol and interfaces have been added as a new feature.

Dedicated OpenFlow switches. A dedicated OpenFlow Switch is a dumb datapath element that forwards packets between ports, as defined by a remote control process. Figure 1 shows an example of an OpenFlow Switch.

In this context, flows are broadly defined, and are limited only by the capabilities of the particular implementation of the Flow Table. For example, a flow could be a TCP connection, or all packets from a particular MAC address or IP address, or all packets with the same VLAN tag, or all packets from the same switch port. For experiments involving non-IPv4 packets, a flow could be defined as all packets matching a specific (but non-standard) header.

Each flow-entry has a simple action associated with it; the three basic ones (that all dedicated OpenFlow switches must support) are:

1. Forward this flow's packets to a given port (or ports). This allows packets to be routed through the network. In most switches this is expected to take place at line-rate.

2. Encapsulate and forward this flow's packets to a controller. Packet is delivered to Secure Channel, where it is encapsulated and sent to a controller. Typically used for the first packet in a new flow, so a controller can decide if the flow should be added to the Flow Table. Or in some experiments, it could be used to forward all packets to a controller for processing.

3. Drop this flow's packets. Can be used for security, to curb denial of service attacks, or to reduce spurious broadcast discovery traffic from end-hosts.

An entry in the Flow-Table has three fields: (1) A packet header that defines the flow, (2) The action, which defines how the packets should be processed, and (3) Statistics, which keep track of the number of packets and bytes for each flow, and the time since the last packet matched the flow (to help with the removal of inactive flows).

In the first generation "Type 0" switches, the flow header is a 10-tuple shown in Table 1. A TCP flow could be specified by all ten fields, whereas an IP flow might not include the transport ports in its definition. Each header field can be a wildcard to allow for aggregation of flows, such as flows in which only the VLAN ID is defined would apply to all traffic on a particular VLAN.

The detailed requirements of an OpenFlow Switch are defined by the *OpenFlow Switch Specification* [6].

OpenFlow-enabled switches. Some commercial switches, routers and access points will be enhanced with

In	VLAN	Ethernet			IP			TCP	
Port	ID	SA	DA	Type	SA	DA	Proto	Src	Dst

Table 1: The header fields matched in a "Type 0" OpenFlow switch.

the OpenFlow feature by adding the Flow Table, Secure Channel and OpenFlow Protocol (we list some examples in Section 5). Typically, the Flow Table will re-use existing hardware, such as a TCAM; the Secure Channel and Protocol will be ported to run on the switch's operating system. Figure 2 shows a network of OpenFlow-enabled commercial switches and access points. In this example, all the Flow Tables are managed by the same controller; the OpenFlow Protocol allows a switch to be controlled by two or more controllers for increased performance or robustness.

Our goal is to enable experiments to take place in an existing production network alongside regular traffic and applications. Therefore, to win the confidence of network administrators, OpenFlow-enabled switches must isolate experimental traffic (processed by the Flow Table) from production traffic that is to be processed by the normal Layer 2 and Layer 3 pipeline of the switch. There are two ways to achieve this separation. One is to add a *fourth action*:

4. Forward this flow's packets through the switch's normal processing pipeline.

The other is to define separate sets of VLANs for experimental and production traffic. Both approaches allow normal production traffic that isn't part of an experiment to be processed in the usual way by the switch. All OpenFlow-enabled switches are required to support one approach or the other; some will support both.

Additional features. If a switch supports the header formats and the four basic actions mentioned above (and detailed in the *OpenFlow Switch Specification*), then we call it a "Type 0" switch. We expect that many switches will support additional actions, for example to rewrite portions of the packet header (e.g., for NAT, or to obfuscate addresses on intermediate links), and to map packets to a priority class. Likewise, some Flow Tables will be able to match on arbitrary fields in the packet header, enabling experiments with new non-IP protocols. As a particular set of features emerges, we will define a "Type 1" switch.

Controllers. A controller adds and removes flow-entries from the Flow Table on behalf of experiments. For example, a static controller might be a simple application running on a PC to statically establish flows to interconnect a set of test computers for the duration of an experiment. In this case the flows resemble VLANs in current networks—providing a simple mechanism to isolate experimental traffic from the production network. Viewed this way, OpenFlow is a generalization of VLANs.

One can also imagine more sophisticated controllers that dynamically add/remove flows as an experiment progresses. In one usage model, a researcher might control the complete network of OpenFlow Switches and be free to decide how all flows are processed. A more sophisticated controller might support multiple researchers, each with different accounts and permissions, enabling them to run multiple independent experiments on different sets of flows. Flows identified

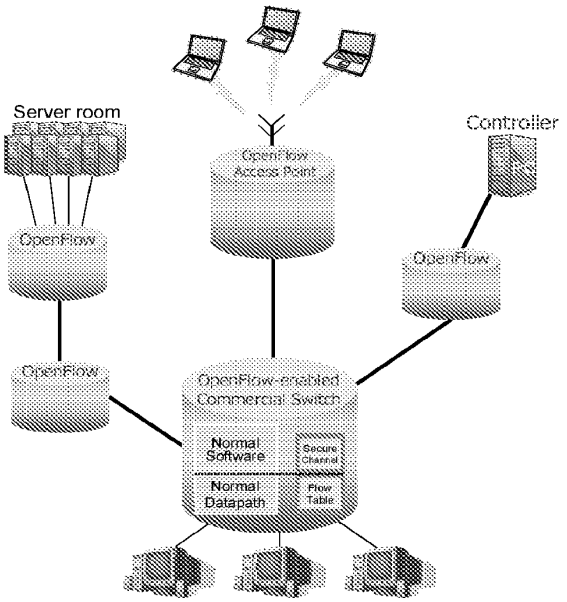


Figure 2: Example of a network of OpenFlow-enabled commercial switches and routers.

as under the control of a particular researcher (e.g., by a policy table running in a controller) could be delivered to a researcher's user-level control program which then decides if a new flow-entry should be added to the network of switches.

3. USING OPENFLOW

As a simple example of how an OpenFlow Switch might be used imagine that Amy (a researcher) invented Amy-OSPF as a new routing protocol to replace OSPF. She wants to try her protocol in a network of OpenFlow Switches, without changing any end-host software. Amy-OSPF will run in a controller; each time a new application flow starts Amy-OSPF picks a route through a series of OpenFlow Switches, and adds a flow- entry in each switch along the path. In her experiment, Amy decides to use Amy-OSPF for the traffic entering the OpenFlow network from her own desktop PC—so she doesn't disrupt the network for others. To do this, she defines one flow to be all the traffic entering the OpenFlow switch through the switch port her PC is connected to, and adds a flow-entry with the action "Encapsulate and forward all packets to a controller". When her packets reach a controller, her new protocol chooses a route and adds a new flow-entry (for the application flow) to every switch along the chosen path. When subsequent packets arrive at a switch, they are processed quickly (and at line-rate) by the Flow Table.

There are legitimate questions to ask about the performance, reliability and scalability of a controller that dynamically adds and removes flows as an experiment progresses: Can such a centralized controller be fast enough to process new flows and program the Flow Switches? What happens when a controller fails? To some extent these questions were addressed in the context of the Ethane prototype, which used simple flow switches and a central controller [7]. Preliminary results suggested that an Ethane controller based on a low-cost desktop PC could process over 10,000 new flows per second — enough for a large college campus. Of course, the rate at which new flows can be processed will depend on the complexity of the processing required by the researcher's experiment. But it gives us confidence that meaningful experiments can be run. Scalability and redundancy are possible by making a controller (and the experiments) stateless, allowing simple load-balancing over multiple separate devices.

3.1 Experiments in a Production Network

Chances are, Amy is testing her new protocol in a network used by lots of other people. We therefore want the network to have two additional properties:

1. Packets belonging to users other than Amy should be routed using a standard and tested routing protocol running in the switch or router from a "name-brand" vendor.

2. Amy should only be able to add flow entries for her traffic, or for any traffic her network administrator has allowed her to control.

Property 1 is achieved by OpenFlow-enabled switches. In Amy's experiment, the default action for all packets that don't come from Amy's PC could be to forward them through the normal processing pipeline. Amy's own packets would be forwarded directly to the outgoing port, without being processed by the normal pipeline.

Property 2 depends on the controller. The controller should be seen as a platform that enables researchers to implement various experiments, and the restrictions of Property 2 can be achieved with the appropriate use of permissions or other ways to limit the powers of individual researchers to control flow entries. The exact nature of these permission-like mechanisms will depend on how the controller is implemented. We expect that a variety of controllers will emerge. As an example of a concrete realization of a controller, some of the authors are working on a controller called NOX as a follow-on to the Ethane work [8]. A quite different controller might emerge by extending the GENI management software to OpenFlow networks.

3.2 More Examples

As with any experimental platform, the set of experiments will exceed those we can think of up-front — most experiments in OpenFlow networks are yet to be thought of. Here, for illustration, we offer some examples of how OpenFlow-enabled networks could be used to experiment with new network applications and architectures.

Example 1: Network Management and Access Control. We'll use Ethane as our first example [7] as it was the research that inspired OpenFlow. In fact, an OpenFlow Switch can be thought of as a generalization of Ethane's datapath switch. Ethane used a specific implementation of a controller, suited for network management and control, that manages the admittance and routing of flows. The basic idea of Ethane is to allow network managers to define a

network-wide policy in the central controller, which is enforced directly by making admission control decisions for each new flow. A controller checks a new flow against a set of rules, such as "Guests can communicate using HTTP, but only via a web proxy" or "VoIP phones are not allowed to communicate with laptops." A controller associates packets with their senders by managing all the bindings between names and addresses — it essentially takes over DNS, DHCP and authenticates all users when they join, keeping track of which switch port (or access point) they are connected to. One could envisage an extension to Ethane in which a policy dictates that particular flows are sent to a user's process in a controller, hence allowing researcher-specific processing to be performed in the network.

Example 2: VLANs. OpenFlow can easily provide users with their own isolated network, just as VLANs do. The simplest approach is to statically declare a set of flows which specify the ports accessible by traffic on a given VLAN ID. Traffic identified as coming from a single user (for example, originating from specific switch ports or MAC addresses) is tagged by the switches (via an action) with the appropriate VLAN ID.

A more dynamic approach might use a controller to manage authentication of users and use the knowledge of the users' locations for tagging traffic at runtime.

Example 3: Mobile wireless VOIP clients. For this example consider an experiment of a new call- handoff mechanism for WiFi-enabled phones. In the experiment VOIP clients establish a new connection over the OpenFlow-enabled network. A controller is implemented to track the location of clients, re-routing connections — by reprogramming the Flow Tables — as users move through the network, allowing seamless handoff from one access point to another.

Example 4: A non-IP network. So far, our examples have assumed an IP network, but OpenFlow doesn't require packets to be of any one format — so long as the Flow Table is able to match on the packet header. This would allow experiments using new naming, addressing and routing schemes. There are several ways an OpenFlow-enabled switch can support non-IP traffic. For example, flows could be identified using their Ethernet header (MAC src and dst addresses), a new EtherType value, or at the IP level, by a new IP Version number. More generally, we hope that future switches will allow a controller to create a generic mask (offset + value + mask), allowing packets to be processed in a researcher-specified way.

Example 5: Processing packets rather than flows. The examples above are for experiments involving flows — where a controller makes decisions when the flow starts. There are, of course, interesting experiments to be performed that require every packet to be processed. For example, an intrusion detection system that inspects every packet, an explicit congestion control mechanism, or when modifying the contents of packets, such as when converting packets from one protocol format to another.

There are two basic ways to process packets in an OpenFlow-enabled network. First, and simplest, is to force all of a flow's packets to pass through a controller. To do this, a controller doesn't add a new flow entry into the Flow Switch — it just allows the switch to default to forwarding every packet to a controller. This has the advantage of flexibility, at the cost of performance. It might provide a useful way to test the functionality of a new protocol, but is unlikely to be of much interest for deployment in a large network.

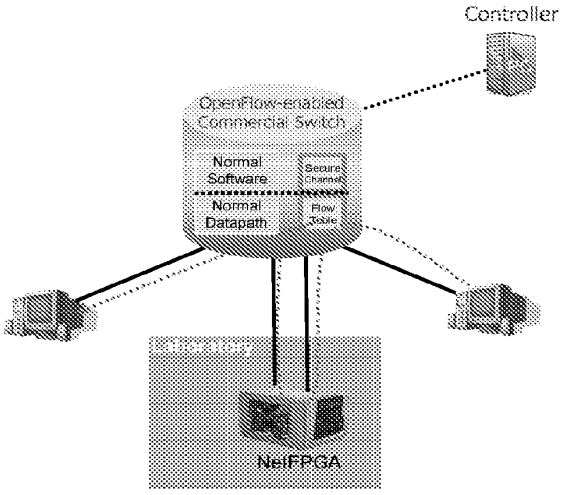


Figure 3: Example of processing packets through an external line-rate packet-processing device, such as a programmable NetFPGA router.

The second way to process packets is to route them to a programmable switch that does packet processing — for example, a NetFPGA-based programmable router. The advantage is that the packets can be processed at line-rate in a user-definable way; Figure 3 shows an example of how this could be done, in which the OpenFlow-enabled switch operates essentially as a patch-panel to allow the packets to reach the NetFPGA. In some cases, the NetFPGA board (a PCI board that plugs into a Linux PC) might be placed in the wiring closet alongside the OpenFlow-enabled switch, or (more likely) in a laboratory.

4. THE OPENFLOW CONSORTIUM

The OpenFlow Consortium aims to popularize OpenFlow and maintain the *OpenFlow Switch Specification*. The Consortium is a group of researchers and network administrators at universities and colleges who believe their research mission will be enhanced if OpenFlow-enabled switches are installed in their network.

Membership is open and free for anyone at a school, college, university, or government agency worldwide. The OpenFlow Consortium welcomes individual members who are not employed by companies that manufacture or sell Ethernet switches, routers or wireless access points (because we want to keep the consortium free of vendor influence). To join, send email to *join@OpenFlowSwitch.org*.

The Consortium web-site [1] contains the OpenFlow Switch Specification, a list of consortium members, and reference implementations of OpenFlow switches.

[1] http://www.OpenFlowSwitch.org

Licensing Model: The OpenFlow Switch Specification is free for all commercial and non-commercial use. (The exact wording is on the web-site.) Commercial switches and routers claiming to be "OpenFlow-enabled" must conform to the requirements of an OpenFlow Type 0 Switch, as defined in the OpenFlow Switch Specification. OpenFlow is a trademark of Stanford University, and will be protected on behalf of the Consortium.

5. DEPLOYING OPENFLOW SWITCHES

We believe there is an interesting market opportunity for network equipment vendors to sell OpenFlow-enabled switches to the research community. Every building in thousands of colleges and universities contains wiring closets with Ethernet switches and routers, and with wireless access points spread across campus.

We are actively working with several switch and router manufacturers who are adding the OpenFlow feature to their products by implementing a Flow Table in existing hardware; i.e. no hardware change is needed. The switches run the Secure Channel software on their existing processor.

We have found network equipment vendors to be very open to the idea of adding the OpenFlow feature. Most vendors would like to support the research community without having to expose the internal workings of their products.

We are deploying large OpenFlow networks in the Computer Science and Electrical Engineering departments at Stanford University. The networks in two buildings will be replaced by switches running OpenFlow. Eventually, all traffic will run over the OpenFlow network, with production traffic and experimental traffic being isolated on different VLANs under the control of network administrators. Researchers will control their own traffic, and be able to add/remove flow-entries.

We also expect many different OpenFlow Switches to be developed by the research community. The OpenFlow website contains "Type 0" reference designs for several different platforms: Linux (software), OpenWRT (software, for access points), and NetFPGA (hardware, 4-ports of 1GE). As more reference designs are created by the community we will post them. We encourage developers to test their switches against the reference designs.

All reference implementations of OpenFlow switches posted on the web site will be open-source and free for commercial and non-commercial use.[2]

[2]Some platforms may limit the license terms of software running on them. For example, a reference implementation on Linux may be limited by the Linux GPL.

6. CONCLUSION

We believe that OpenFlow is a pragmatic compromise that allows researchers to run experiments on heterogeneous switches and routers in a uniform way, without the need for vendors to expose the internal workings of their products, or researchers to write vendor-specific control software.

If we are successful in deploying OpenFlow networks in our campusses, we hope that OpenFlow will gradually catch-on in other universities, increasing the number of networks that support experiments. We hope that a new generation of control software emerges, allowing researchers to re-use controllers and experiments, and build on the work of others. And over time, we hope that the islands of OpenFlow networks at different universities will be interconnected by tunnels and overlay networks, and perhaps by new OpenFlow networks running in the backbone networks that connect universities to each other.

7. REFERENCES

[1] Global Environment for Network Innovations. Web site http://geni.net.

[2] Mark Handley Orion Hodson Eddie Kohler. "XORP: An Open Platform for Network Research," *ACM SIGCOMM Hot Topics in Networking, 2002.*

[3] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. Frans Kaashoek. "The Click modular router," *ACM Transactions on Computer Systems 18(3), August 2000, pages 263-297.*

[4] J. Turner, P. Crowley, J. Dehart, A. Freestone, B. Heller, F. Kuhms, S. Kumar, J. Lockwood, J. Lu, M.Wilson, C. Wiseman, D. Zar. "Supercharging PlanetLab - High Performance, Multi-Application, Overlay Network Platform," *ACM SIGCOMM '07, August 2007, Kyoto, Japan.*

[5] NetFPGA: Programmable Networking Hardware. Web site http://netfpga.org.

[6] The OpenFlow Switch Specification. Available at http://OpenFlowSwitch.org.

[7] Martin Casado, Michael J. Freedman, Justin Pettit, Jianying Luo, Nick McKeown, Scott Shenker. "Ethane: Taking Control of the Enterprise," *ACM SIGCOMM '07, August 2007, Kyoto, Japan.*

[8] Natasha Gude, Teemu Koponen, Justin Pettit, Ben Pfaff, Martin Casadao, Nick McKeown, Scott Shenker, "NOX: Towards an Operating System for Networks," *In submission. Also: http://nicira.com/docs/nox-nodis.pdf.*



A SHORT INTRODUCTION TO CLOUD PLATFORMS

AN ENTERPRISE-ORIENTED VIEW

DAVID CHAPPELL

AUGUST 2008

SPONSORED BY MICROSOFT CORPORATION

COPYRIGHT © 2008 CHAPPELL & ASSOCIATES

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0061

CONTENTS

Defining Terms: What is a Cloud Platform? 3

Cloud Platforms in Context: Three Kinds of Cloud Services 3

A General Model for Application Platforms 4

From On-Premises Platforms to Cloud Platforms 7

Examining Cloud Platforms 7

Cloud Foundation 8

Operating System 8

Local Support 8

Cloud Infrastructure Services 9

Storage 9

Integration 10

Identity 11

Cloud Application Services 11

SaaS Application Services 11

Search 12

Mapping 12

Other Application Services 12

Conclusion 13

About the Author 13

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0062

DEFINING TERMS: WHAT IS A CLOUD PLATFORM?

The coming shift to cloud computing is a major change in our industry. One of the most important parts of that shift is the advent of *cloud platforms*. As its name suggests, this kind of platform lets developers write applications that run in the cloud, or use services provided from the cloud, or both. Different names are used for this kind of platform today, including *on-demand platform* and *platform as a service (PaaS)*. Whatever it's called, this new way of supporting applications has great potential.

To see why, think about how application platforms are used today. When a development team creates an on-premises application (i.e., one that will run within an organization), much of what that application needs already exists. An operating system provides basic support for executing the application, interacting with storage, and more, while other computers in the environment offer services such as remote storage. If the creators of every on-premises application first had to build all of these basics, we'd have many fewer applications today.

Similarly, if every development team that wishes to create a cloud application must first build its own cloud platform, we won't see many cloud applications. Fortunately, vendors are rising to this challenge, and a number of cloud platform technologies are available today. The goal of this overview is to categorize and briefly describe those technologies as they're seen by someone who creates enterprise applications.

CLOUD PLATFORMS IN CONTEXT: THREE KINDS OF CLOUD SERVICES

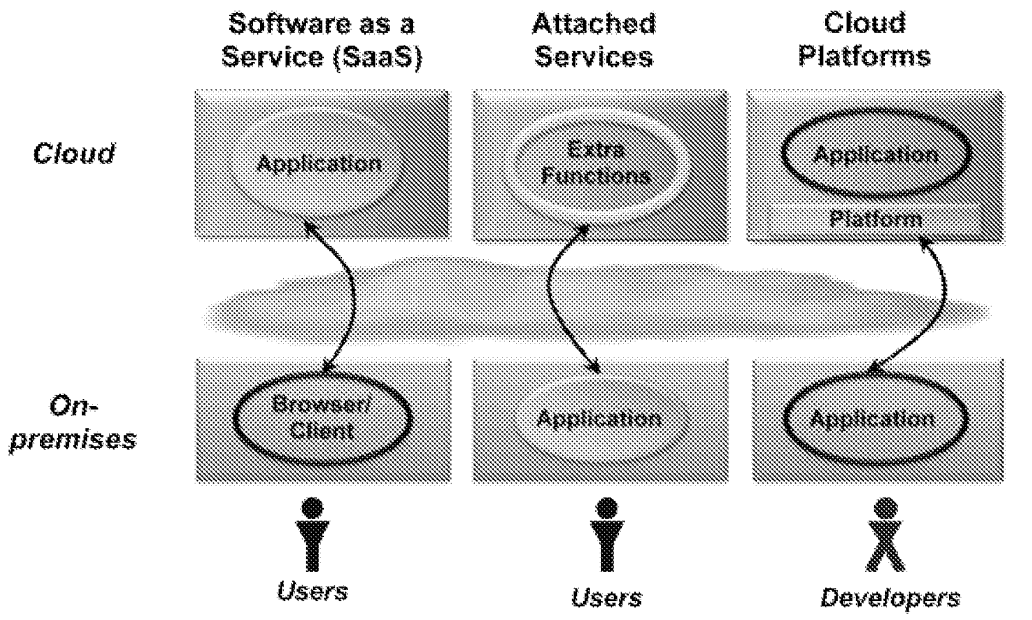


Figure 1: Cloud services can be grouped into three broad categories.

To get a grip on cloud platforms, it's useful to start by looking at cloud services in general. As Figure 1 shows, services in the cloud can be grouped into three broad categories. Those categories are:

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0063

- Software as a service (SaaS): A SaaS application runs entirely in the cloud (that is, on servers at an Internet-accessible service provider). The on-premises client is typically a browser or some other simple client. The most well-known example of a SaaS application today is probably Salesforce.com, but many, many others are also available.

- Attached services: Every on-premises application provides useful functions on its own. An application can sometimes enhance these by accessing application-specific services provided in the cloud. Because these services are usable only by this particular application, they can be thought of as attached to it. One popular consumer example of this is Apple's iTunes: The desktop application is useful for playing music and more, while an attached service allows buying new audio and video content. Microsoft's Exchange Hosted Services provides an enterprise example, adding cloud-based spam filtering, archiving, and other services to an on-premises Exchange server.

- Cloud platforms: A cloud platform provides cloud-based services for creating applications. Rather than building their own custom foundation, for example, the creators of a new SaaS application could instead build on a cloud platform. As Figure 1 shows, the direct users of a cloud platform are developers, not end users.

Understanding cloud platforms requires some agreement on what the word "platform" means in this context. One broad way to think about it is to view a platform as any software that provides developer-accessible services for creating applications. The next section looks at this idea in a bit more detail.

A GENERAL MODEL FOR APPLICATION PLATFORMS

Our experience with application platforms today comes mostly from on-premises platforms. A useful way to think about cloud platforms is to see how the services an application developer relies on in the on-premises environment translate to the cloud. To help do this, Figure 2 shows a general model that can be applied to both worlds.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0064

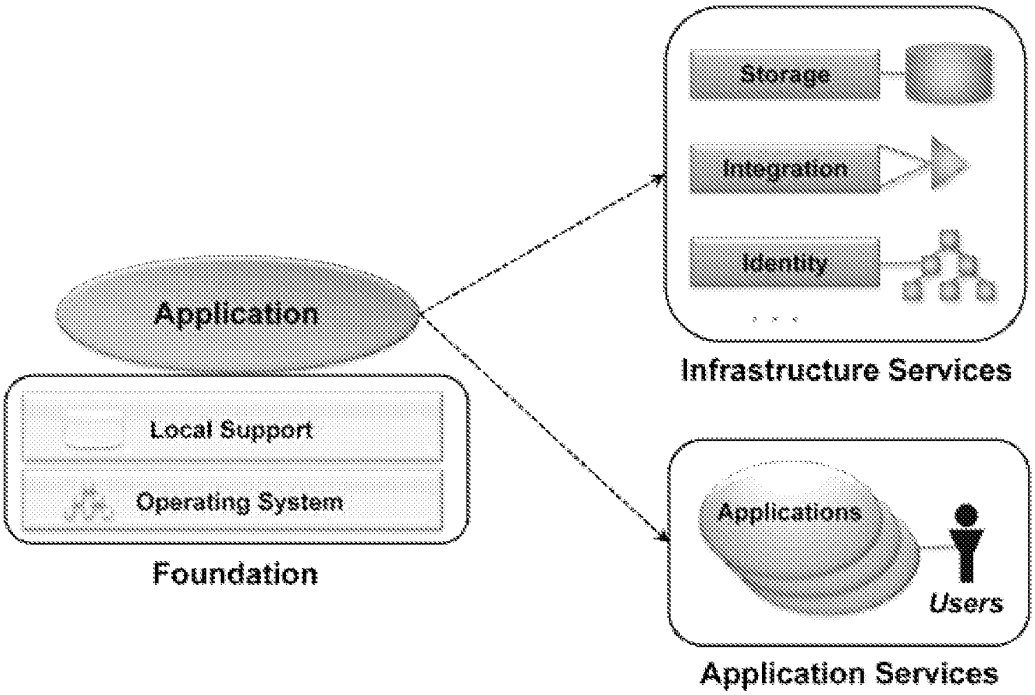


Figure 2: A modern application platform can be viewed as having three parts.

Whether it's on-premises or in the cloud, an application platform can be thought of as comprising three parts:

- A *foundation*: Nearly every application uses some platform software on the machine it runs on. This typically includes various support functions, such as standard libraries and storage, and a base operating system.
- A group of *infrastructure services*: In a modern distributed environment, applications frequently use basic services provided on other computers. It's common to provide remote storage, for example, integration services, an identity service, and more.
- A set of *application services*: As more and more applications become service-oriented, the functions they offer become accessible to new applications. Even though these applications exist primarily to provide services to end users, this also makes them part of the application platform. (It might seem odd to think of other applications as part of the platform, but in a service-oriented world, they certainly are.)

And while they're not shown in Figure 2, development tools are another important part of this story. Modern tools can help developers build applications using all three parts of an application platform.

To make this abstract model more concrete, think about how it fits with today's most popular on-premises platforms. The on-premises foundation looks like this:

- Operating system: The dominant choices are Windows, Linux, and other versions of Unix.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0065

- Local support: Different technologies are used depending on the style of application. The .NET Framework and Java EE application servers provide general support for Web applications and more, for instance, while other technologies target specific kinds of applications. For example, Microsoft's Dynamics CRM product includes a platform designed for creating a particular type of business application. Similarly, different kinds of storage are used for different purposes. Raw byte storage is provided by the file systems in Windows, Linux, and other operating systems, while more structured storage is provided by a range of database technologies, including the Oracle DBMS, MySQL, Microsoft SQL Server, and IBM DB2.

For on-premises infrastructure services, typical examples include the following:

- Storage: Like storage in the foundation, infrastructure storage comes in various styles. A remote file system might provide simple byte-oriented storage, while a Microsoft SharePoint document library provides a more structured remote storage service. Applications can also access a database system remotely, allowing access to another kind of structured storage.

- Integration: Connecting applications within an organization usually depends on a remote service provided by some integration product. A message queue is a simple example of this, while more complex scenarios use products such as IBM's WebSphere Process Server, Microsoft's BizTalk Server, and others.

- Identity: Providing identity information is a fundamental requirement for most distributed applications. Common on-premises technologies that address this include Microsoft's Active Directory and other LDAP servers.

On-premises application services, the third category shown in Figure 2, vary widely across different organizations. The reason for this is simple: Different organizations use different applications, which in turn expose diverse services. One way to think about these applications in the on-premises platform is to divide them into two broad categories:

- Packaged applications: This includes business software such as SAP, Oracle Applications, and Microsoft Dynamics, along with a myriad of other off-the-shelf products. While not all packaged applications expose services to other applications, more and more of them do.

- Custom applications: Many organizations have a large investment in custom software. As these applications increasingly expose their functionality through services, they become part of the on-premises application platform.

When it's described like this, the on-premises application platform can seem quite complex. The truth, though, is that this platform has evolved over time. In the early days of computing, the application platform consisted of nothing more than an on-premises foundation. (Think of MVS and IMS on an IBM mainframe, for example.) In the 1980s and 1990s, as distributed computing spread, on-premises infrastructure services were added, with remote storage, integration, and identity becoming common. Today, with the advent of service-oriented applications, on-premises application services have become part of the platform. The next step in this evolution is clear: providing cloud versions of all three.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0066

FROM ON-PREMISES PLATFORMS TO CLOUD PLATFORMS

Along with describing on-premises platforms, the general model just described can also be used to think about cloud platforms. And since on-premises and cloud platforms can be used together, it's important to understand how the two work in concert. Figure 3 illustrates this new world.

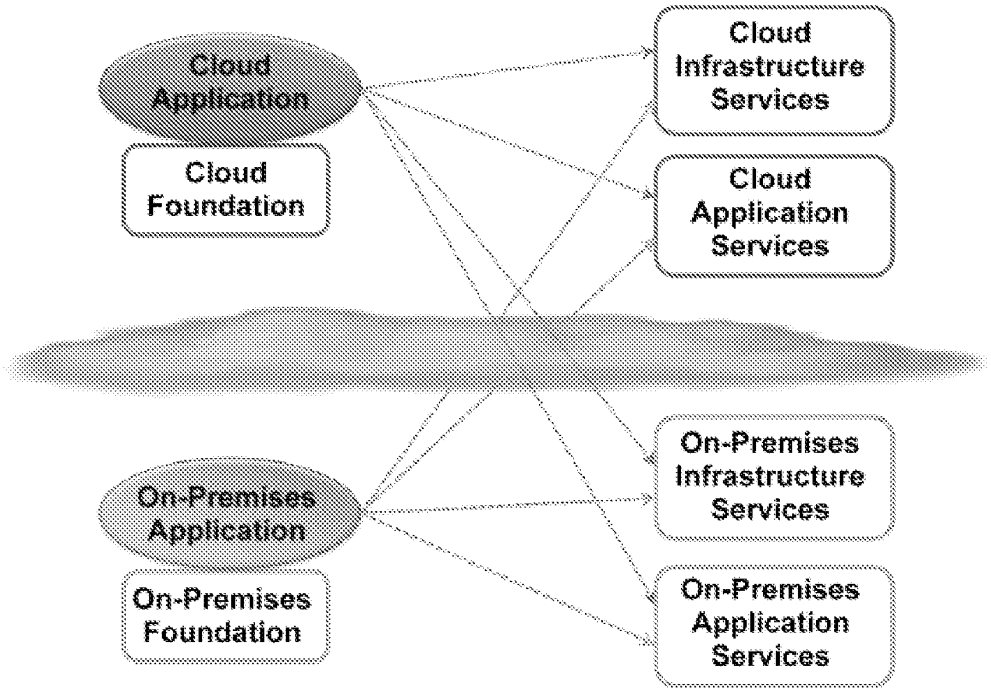


Figure 3: On-premises platforms and cloud platforms can be viewed in similar ways, and they can also be used together.

As the figure shows, a cloud application can be built on a cloud foundation, just as an on-premises application is built on an on-premises foundation. Both kinds of applications can access infrastructure and application services provided on-premises and in the cloud. Just as on-premises platforms support today's applications, cloud platforms provide services for the applications we're likely to build tomorrow.

EXAMINING CLOUD PLATFORMS

Understanding cloud platforms means looking at each of their parts: the cloud foundation, cloud infrastructure services, and cloud application services. This section walks through these three areas, using some of today's most visible cloud platform technologies as examples.

Before we begin, one important note: While it's useful to look at on-premises platforms and cloud platforms through the same lens, the two aren't identical. When platform functions move into the cloud, they sometimes change in significant ways. For example, on-premises platforms are designed to support (at most) enterprise-scale applications. Applications that run in the cloud, by contrast, can potentially operate at Internet scale, which requires handling many more simultaneous users than any enterprise application. While the same kinds of platform functions might be needed in both cases, achieving this high

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0067

scalability can force a cloud platform to provide them in a quite different way. In what follows, expect to see some differences from the on-premises world.

CLOUD FOUNDATION

Like their on-premises cousins, cloud foundations provide the basic local functions an application needs. These can include an underlying operating system and local support. Yet how cloud platforms provide these functions differs from what we're used to, as this section shows.

Operating System

From a platform point of view, an operating system provides a set of basic interfaces for applications to use. By far the most well-known example of an operating system in the cloud today is Amazon's Elastic Compute Cloud (EC2). EC2 provides customer-specific Linux instances running in virtual machines (VMs). From a technical perspective, it might be more accurate to think of EC2 as a platform for VMs rather than operating systems. Still, a developer sees an operating system interface, and so viewing it in this light makes more sense here.

Each development team is free to use whatever local support it likes in this VM—Amazon doesn't care. The creators of one application might choose a Java EE app server and MySQL, for example, while another group might go with Ruby on Rails. EC2 customers are even free to create many Linux instances, then distribute large workloads across them in parallel, such as for scientific applications. While the service EC2 provides is quite basic, it's also very general, and so it can be used in many different ways.

Local Support

In an on-premises platform (and in EC2), a developer can mix and match parts of the foundation as she sees fit. Choosing to use the .NET Framework on Windows doesn't mandate using a particular database, for example. Similarly, an on-premises application using the .NET Framework is free to access the underlying Windows operating system, as is an application built on a Java EE server.

The local support functions in today's leading cloud foundations don't work this way. Instead, a cloud local support technology typically includes its own storage, and it hides whatever the underlying operating system might be. A developer choosing to build on a particular local support option must accept the limitations it imposes.

There are good reasons for these limitations, of course. One of the things that makes cloud computing so attractive is its potential for scalability, but to make an application built on a cloud foundation handle Internet-size loads requires limiting it in some ways. By making the local support functions more specialized, a cloud platform provider has more freedom to optimize the application environment. Accordingly, each set of local support functions in cloud foundations today focuses on supporting a particular kind of application.

For example, Google's AppEngine provides local support for running Python Web applications. Along with a standard Python runtime, AppEngine also includes a hierarchical data store with its own query language. Another example of a cloud platform providing local support is Force.com, offered by Salesforce.com. Rather than targeting general Web applications, however, Force.com is aimed at creating data-oriented

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0068

business applications. Toward this end, it provides its own focused support for data storage. And rather than adopt an existing programming language, this platform's creators invented their own, a language called Apex.

Microsoft also provides local support for applications in the cloud as part of its CRM Live offering. Based on the Dynamics CRM platform mentioned earlier, this technology targets data-oriented business applications, much like Force.com. And like both Force.com and AppEngine, it includes both run-time application support and a data store. Microsoft has also talked about its plans to go further in this area, with a platform that will support standard .NET development languages and tools. The intent, Microsoft says, is to allow portability of both applications and developer skills between the company's on-premises foundation and its cloud foundation.

CLOUD INFRASTRUCTURE SERVICES

Whether they run on-premises or in the cloud, some applications don't need anything beyond a foundation. Still, many can benefit from distributed storage, common identity, and other infrastructure services. We're accustomed to having these services provided on-premises today, but analogous services are also provided in the cloud.

As Figure 3 showed, cloud infrastructure services can be accessed by applications running on either an on-premises foundation or a cloud foundation. Initially, the most common users of cloud infrastructure services will be on-premises, because there aren't yet many applications built on a cloud foundation. Over time, expect this to change, as more and more cloud-based applications also use cloud infrastructure services.

Storage

Applications commonly use some kind of local storage, which is why storage is part of both on-premises and cloud foundations. Remote storage is also useful, however, as the popularity of this service in the on-premises world shows. Accordingly, it's reasonable to expect that providing a storage service in the cloud will be attractive for many applications.

As with on-premises platforms, remote storage in the cloud comes in different styles. For example, Amazon's Simple Storage Service (S3) provides basic unstructured remote storage. The model it exposes to developers is straightforward: *objects*, which are just bunches of bytes, are stored in *buckets*. Applications can create, read, and delete objects and buckets. Objects can't be updated, however—they can only be entirely replaced. This is another example of how platform services must change to support Internet-scale usage, something that Amazon is clearly focused on. This simple but limited storage service is much easier to make scalable than a more fully featured offering would be. The trade-off is clear: Application developers get cheap storage in the cloud, but they might need to do more work in their applications to use it effectively.

Another approach to cloud storage is to support more structured data. In Microsoft's SQL Server Data Services (SSDS), for example, a *container* includes one or more *entities*, each of which holds some number of *properties*, as shown in Figure 4. An application can issue queries against a container's data with operators such as ==, !=, <, >, AND, OR, and NOT.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0069

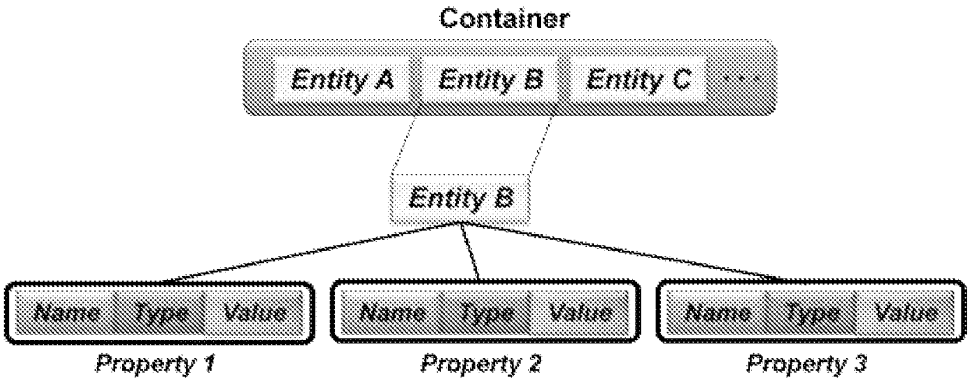


Figure 4: In SQL Server Data Services, a container holds entities with properties.

It's important to note that this isn't a relational database, and the query language isn't SQL. Once again, we're seeing an illustration of how application platform technologies change when they're moved into the cloud. This simpler approach is easier to use than a relational database—there's no need to define a schema up front—and it's also easier to make scalable.

Amazon's SimpleDB provides one more example of the value of structured storage in the cloud. The way SimpleDB organizes information is similar to SSDS—it's a hierarchy of *domains*, *items*, and *values*—and it also provides a non-SQL query language. Like SSDS, no up-front schema definition is required, and so the approach provides a mix of flexibility and scalability.

Integration

Is there any application left that doesn't talk to at least one of its fellows? Connecting applications has become a staple of computing, and vendors have provided a plethora of on-premises infrastructure services to do it. These range from relatively simple technologies like message queues to quite complex integration servers.

As integration services move into the cloud, a range of technologies is also appearing. For example, Amazon's Simple Queue Service (SQS) provides just what its name suggests: a straightforward way for applications to exchange messages via queues in the cloud. Yet SQS once again illustrates what happens when a familiar on-premises service is recast as a cloud service. Because SQS replicates messages across multiple queues, an application reading from a queue isn't guaranteed to see all messages from all queues on a particular read request. SQS also doesn't promise in-order, exactly-once delivery. These simplifications let Amazon make SQS more scalable, but they also mean that developers must use SQS differently from an on-premises message queuing technology.

BizTalk Services provides another example of cloud-based integration. Rather than using message queuing, BizTalk Services implements a relay service in the cloud that lets applications communicate through firewalls. Cloud-based integration, such as connecting applications in different organizations, typically requires traversing firewalls, and so solving this problem is important. BizTalk Services also provides simple workflow support along with a way for an application to register the services it exposes, then let those services be invoked by any other application that has permission to do so.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0070

Going forward, expect to see more integration services offered in the cloud. Given the importance of integration as an on-premises service, it shouldn't be surprising to see its functions become part of the cloud infrastructure.

Identity

Whether an application runs on-premises or in the cloud, it typically needs to know something about its users. Toward this end, the application commonly demands that each user provides a digital identity, a set of bytes that describes that user. Based on what these bytes contain and how they're verified, the application can determine things such as who this user is and what they're allowed to do.

Many on-premises applications today rely on an on-premises infrastructure service, such as Active Directory, to provide this identity information. When a user accesses a cloud application, however, or an on-premises application accesses a cloud service, an on-premises identity usually won't work. And what about an application built on a cloud foundation? Where does it get its identity information?

An identity service in the cloud can address these issues. Because it provides a digital identity that can be used by people, by on-premises applications, and by cloud applications, a cloud identity service can be applied in many different scenarios. In fact, one indication of the importance of this kind of identity service is the number of cloud identity services available today. Accessing Amazon cloud services such as EC2 or S3 requires presenting an Amazon-defined identity, for instance, while using Google AppEngine requires a Google account. Microsoft provides Windows Live ID, which can be used for Microsoft applications and others, while BizTalk Services also offers its own identity service, which can be federated with others. Developers don't have complete freedom—cloud platforms are frequently tied to a particular identity provider—but the need for identity as a cloud service is clear.

CLOUD APPLICATION SERVICES

What's the difference between an application service and an infrastructure service? To answer this question, think first about the obvious distinction between applications and infrastructure: Applications are designed to be used by people, while infrastructure is designed to be used by applications. It's also fair to say that infrastructure usually provides a general, relatively low-level service, while applications provide more specific, higher-level services. An infrastructure service solves a broad problem faced by many different kinds of applications, while an application service solves a more targeted problem. And just as it's possible to identify different kinds of infrastructure services, it's also possible to distinguish different categories of application services, as this section illustrates.

SaaS Application Services

Users in most enterprises today rely on both purchased and home-grown applications. As these applications expose their services to remote software, they become part of the on-premises platform. Similarly, SaaS applications today frequently expose services that can be accessed by on-premises applications or by other cloud applications. Salesforce.com's CRM application, for example, makes available a variety of services that can be used to integrate its functions with on-premises applications. As organizations begin to create their own SaaS applications running on a cloud foundation, those applications will also expose services. Just as packaged and custom on-premises applications today are

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0071

part of the on-premises platform, the services exposed by packaged and custom SaaS applications are becoming part of the cloud platform.

Search

Services exposed by SaaS applications are useful, but they're not the whole story. Other kinds of cloud application services are also important. Think, for example, of search engines such as Google and Live Search. Along with their obvious value to people, why can't they also offer cloud application services?

The answer, of course, is that they can. Microsoft's Live Search, for example, exposes services that allow on-premises and cloud applications to submit searches and get results back. Suppose a company that provided a database of legal information wanted to let customers search both its own data and the Web in a single request. They could accomplish this by creating an on-premises application that both searched their proprietary data and, via the Live Search application service, the entire Web. It's fair to say that not many applications are likely to need this kind of service, but that's one reason why it's most accurate to think of search as an application service rather than an infrastructure service.

Mapping

Many Web applications today display maps. Hotel Web sites plot their locations, retailers provide store locators, and more. The people who create these applications probably don't have the time, interest, or budget to create their own mapping database. Yet enough applications need this function to justify creating a cloud application service that provides it.

This is exactly what's done by mapping services such as Google Maps and Microsoft's Virtual Earth. Both provide cloud-based services that application developers can use to embed maps in Web pages and more. And as with search, these mapping services are adjuncts to existing Web sites that target users directly, i.e., they're cloud application services.

Other Application Services

Many other application services are available today. In fact, almost any Web site can expose its functionality as a cloud service for developers to use. Photo-sharing sites such as Google's Picasa and Microsoft's Windows Live Photo Gallery do this, for example, as do online contacts applications such as Google Contacts and Microsoft's Windows Live Contacts. One big motivation for exposing services is to make it easier to create mash-ups that exploit the functions of diverse Web applications.

Vendors sometimes group cloud application services together under a common umbrella. The services for accessing information in Google Contacts, Picasa, and more are all part of the Google Data APIs, for instance. Similarly, Microsoft groups several of its services together under the Live Platform brand, including Live Search, Virtual Earth, Windows Live Contacts, Windows Live ID, an Alerts service, a specialized storage service called Application-Based Storage, and several more.

The line between cloud infrastructure services and cloud application services can sometimes be fuzzy. General cloud storage services such as S3 and SSDS are clearly infrastructure, for example, as are cloud identity services. A mapping service such as Google Earth is just as clearly application-centric—only certain kinds of apps need it—as is a service like Live Search. But an Alerts service might be considered

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0072

infrastructure, since it's more generally useful, and Windows Live ID is definitely infrastructure, even though Microsoft views both services as part of its Live Platform.

Cloud platforms are a relatively new area, and so it shouldn't be surprising that defining a firm taxonomy is challenging. However you choose to view them, it's clear that cloud application services have an important role to play. Knowing what's available in the cloud should be a core competency today for everyone who designs and builds software.

CONCLUSION

A new kind of application platform doesn't come along very often. But when a successful platform innovation does appear, it has an enormous impact. Think of the way personal computers and servers shook up the world of mainframes and minicomputers, for example, or how the rise of platforms for N-tier applications changed the way people write software. While the old world doesn't go away, a new approach can quickly become the center of attention for new applications.

Cloud platforms don't yet offer the full spectrum of an on-premises environment. For example, business intelligence as part of the platform isn't common, nor is support for business process management technologies such as full-featured workflow and rules engines. This is all but certain to change, however, as this technology wave continues to roll forward.

Cloud platforms aren't yet at the center of most people's attention. The odds are good, though, that this won't be true five years from now. The attractions of cloud-based computing, including scalability and lower costs, are very real. If you work in application development, whether for a software vendor or an end user, expect the cloud to play an increasing role in your future. The next generation of application platforms is here.

ABOUT THE AUTHOR

David Chappell is Principal of Chappell & Associates (www.davidchappell.com) in San Francisco, California. Through his speaking, writing, and consulting, he helps software professionals around the world understand, use, and make better decisions about new technology.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0073

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	Distributed Software Defined Networking

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [■] The attached application, or
- [ ] United States application or PCT international application number _____

filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Vivek Ramanna Date (Optional): June 2, 2014

Signature: Vivek

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	Distributed Software Defined Networking

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [x] The attached application, or
- [ ] United States application or PCT international application number _____ filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Pouya Taaghol Date (Optional): June 2, 2014

Signature: _____

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Electronic Patent Application Fee Transmittal

Application Number:	
Filing Date:	
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Filer:	Christopher James Horgan
Attorney Docket Number:	TAAG001CONT1

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
UTILITY FILING FEE (ELECTRONIC FILING)	4011	1	70	70
UTILITY SEARCH FEE	2111	1	300	300
UTILITY EXAMINATION FEE	2311	1	360	360
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**730**

Electronic Acknowledgement Receipt

EFS ID:	31179296
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	09-DEC-2017
Filing Date:	
Time Stamp:	00:23:05
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Specification	TAAG001CONT1_SPECIFICATION.pdf	879587 9a28f4d793d2f110c3d32126fbd49fd58868182b	no	61

Warnings:

Information:					
2	Drawings-only black and white line drawings	TAAG001CONT1_DRAWINGS.pdf	1280529 f4d0baffdbde2fcfd0a8951ca044e172ce7242d5	no	18
Warnings:					
Information:					
3	Application Data Sheet	TAAG001CONT1_ADS.pdf	1611116 7836eb3ac91dd422c53448b30f8947e101e98301	no	6
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
4	Power of Attorney	TAAG001CONT1_POA.pdf	1016501 d42b0598edb939be83036dfe8849846128130d65	no	2
Warnings:					
Information:					
5	Internet Communications Authorized	TAAG001CONT1_INT_COMM.pdf	78889 47c22e9793cf22126a48998e4d05286a6aba0738	no	2
Warnings:					
Information:					
6	Nonpublication request from applicant	TAAG001CONT1_NON_PUB.pdf	121599 a8f87137addce261fa45569e4c09ac3362fd0a7b	no	1
Warnings:					
Information:					
7	Information Disclosure Statement (IDS) Form (SB08)	TAAG001CONT1_IDS.pdf	322552 b178a91b245dfdb11221b9d250c6aea47787685c	no	3
Warnings:					
Information:					
This is not an USPTO supplied IDS fillable form					
8	Non Patent Literature	TAAG001CONT1_NON_PATENT_LIT.pdf	6046334 bc01a4899a1deee3221d1d554ee47d9c877a45c5	no	64
Warnings:					

Information:					
9	Oath or Declaration filed	DECLARATION___RAMANNA___TAAG001CONT1.pdf	458897 dfcacd8ddd902164c4639f747b7c2fda44517973	no	1
Warnings:					
Information:					
10	Oath or Declaration filed	TAAG001CONT1_DECLARATION_TAAGHOL.pdf	460756 621e3246b02211653607b4f8a4a081a74780e3ac	no	1
Warnings:					
Information:					
11	Fee Worksheet (SB06)	fee-info.pdf	34977 8f948cc513905dacb7a0613f9ca11b566f40d8ca	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	12311737		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

DISTRIBUTED SOFTWARE DEFINED NETWORKING

CROSS-REFERENCE TO RELATED PATENT APPLICATIONS

[0001] The present invention claims priority from U.S. Patent Application Ser. No. 14/295,331, filed June 4, 2014; which claims priority from U.S. Provisional Patent Application Serial No. 61/834,807, filed June 13, 2013; both of these disclosures are herein specifically incorporated by reference in their entirety.

FIELD OF THE INVENTION

[0002] This invention relates to a network architecture that facilitates secure and flexible programmability between a user device and across a network with full lifecycle management of services and infrastructure applications.

BACKGROUND

[0003] Existing network solutions are built on proprietary hardware and software. Network Operators & Information Technology (NOIT) can only configure the proprietary network infrastructure provided by vendors and are unable to add new and customized features or capabilities. As a result, the NOIT can only add new desired features by making such requests to their infrastructure vendors or pursuing standardization processes. But these existing approaches are time consuming and resourceintensive.

SUMMARY OF THE INVENTION

[0004] An aspect of the disclosure herein is a plurality of network elements in a packet system, said network elements comprising: a flexible network device adapted to process data flows in the packet system having at least one fxDeviceApp; a flexible cloud platform having at least one fxCloudApp, wherein the at least one fxDeviceApp and at least one fxCloudApp are in secure communication to form a distributed application; and an application management portal which manages life cycle of the distributed applications.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0081

[0005] Another aspect of the disclosure is a plurality of network elements in a packet system, said network elements comprising: a flexible cloud platform having a plurality of virtual machines, wherein each of the virtual machines has at least one fxCloudApp; and a flexible network device adapted to process data flows in the packet system having at least one fxDeviceApp, wherein the at least one fxDeviceApp and fxCloudApp are in secure communication to form a distributed application.

[0006] Another aspect of the disclosure is a plurality of network elements in a packet system, said network elements comprising: a flexible cloud platform having a plurality of zones, wherein each of the plurality of zones has at least one fxCloudApp; and a flexible network device adapted to process data flows in the packet system having at least one fxDeviceApp, wherein the at least one fxDeviceApp and fxCloudApp are in secure communication to form a distributed application.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Figure 1 demonstrates the network architecture of a 4G system.

[0008] Figure 2 illustrates a network element (or system) upon which the Distributed Software Defined Network operates.

[0009] Figure 3 illustrates a Distributed Software Defined Network in accordance with the present disclosure.

[0010] Figure 4 is an exemplary view of the software layers in the Flexible Operation System (fxOS) Architecture.

[0011] Figure 5 illustrates the FastPath architecture.

[0012] Figure 6 presents a potential software layer of the fxOS co-existing alongside a general purpose OS.

[0013] Figure 7 is an exemplary view of the fxStore Portal.

[0014] Figure 8 demonstrates a shared fxCloud deployment.

[0015] Figure 9 demonstrates a Zoned Deployment.

[0016] Figure 10 demonstrates another deployment option as applied to cellular networks (especially Long Term Evolution (LTE)).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0082

[0017] Figure 11 shows an example of a signaling flow chart for the fxDevice Commissioning Procedure.

[0018] Figure 12 shows an example procedure for Application Provisioning.

[0019] Figure 13 presents the state machine for the application upgrade process.

[0020] Figures 14A-14C are examples of the Distributed Switch Controller during Upgrade.

[0021] Figure 15 shows how fxVF in the fxCloud and fxDevice.

[0022] Figure 16 demonstrates the logical interfaces between the platform resources and applications.

[0023] Figures 17A and 17B demonstrate example procedures for cloud breathing procedures.

DETAILED DESCRIPTION

[0024] The Distributed Software Defined Network (dSDN) disclosed herein is an architecture that enables secure and flexible programmability across a network with full lifecycle management of services and infrastructure applications (fxDeviceApp). The dSDN also harmonizes application deployment across the network independent of the hardware vendor. As a result, the dSDN simplifies the network deployment lifecycle from concept to design to implementation to decommissioning.

DEFINITIONS, ACRONYMS & ABBREVIATIONS

[0025] The following terms, acronyms, abbreviations and descriptions are explained below and are used throughout the detailed description of the dSDN:

TERM	DESCRIPTION
2G	2nd Generation Cellular Technology
3G	3Rd Generation Cellular Technology
4G	4th Generation Cellular Technology
AP	WiFi Access Point
API	Application Programming Interface

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0083

AR	Application Repository
ASIC	Application Specification Integrated Circuit
AuC	Authentication Center
BS	Base Station
BSC	Base Station Controller
BTS	Base Station
CDN	Content Distribution Network
CN	Core Network
CNIE	Core Network Internet Edge
CPU	Central Processing Unit
dAP	Distributed Application Package
dApp	Distributed App
dCP	Distributed Content Service
dNS	Distributed Notification Service
DNS	Domain Name System
DPDK	Data Plane Development Kit
DPI	Deep Packet Inspection
dRS	Distributed Resource Service
DSCP	Differentiated Service Code Point (QoS)
E-UTRAN	Evolved UMTS Terrestrial Radio Access Network
EIR	Equipment Identity Register
eNB	Enhanced Node B (4G)
FIB	Flow Information Base
FQDN	Fully Qualified Domain Name
fxApp	FleXible Application, which has a fxDeviceApp and fxCloudApp component. fxDeviceApp and fxCloudApp of one application represent fxApp.
fxDeviceApp	FleXible Application which runs on fxDevice
fxBS	FleXible Base Station that is a BS that is powered by fxOS (an example of an fxDevice)

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0084

fxCloud	FleXible Cloud that interacts with fxDevice and on other network elements on the northbound Interface
fxCloudApp	FleXible Cloud Application that is the application that runs on fxCloud and can be associated to one or more fxDeviceApp in fxDevice
fxManager	FleXible Manager that manager the fxOS lifecycle (provisioning, usage, de-provisioning)
fxOS	FleXible Operation System which run on the fxDevice as the OS and firmware
fxSDK:	FleXible Software Development Kit
fxSimulator	FleXible Simulator
fxStore	FleXible Store that presents fxApp to the network
administratorsGERAN	GPRS Edge Radio Access Network
GGSN	Gateway GPRS Support Node (used in core network of 2G/3G systems)uhu
GMSC	Gateway Mobile Switching Center
GPRS	General Packet Radio Service
GPS	Global Positioning System
GTP	GPRS Tunneling Protocol
GW	Gateway
HAL	Hardware Abstraction Layer
HD	High Definition
HLR	Home Location Register
HSS	Home Subscriber Server
IAC	Inter Application Communication

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0085

IMS	IP Multimedia Sub-system
IMSI	International Mobile Subscriber Identifier
IMSI	International Mobile Subscriber Identity
IP	Internet Protocol
JVM	Java Virtual Machine
LA	Location Area
LAN	Local Area Network
LSP	Label Switching Protocol
M2M	Machine to Machine
MAC	Medium Access Control
MCP	Multi Chip Packaging
MIMO	Multi Input Multi Output
MME	Mobility Management Entity (used in core network of LTE)
MNO	Mobile Operator Network
MPLS	Multi Protocol Label Switching
MSC	Mobile Switching Center
MSISDN	Mobile Station
NAS	Non-Access Stratum
NB	Node B (3G)
NFC	Near Field Communication
NFV	Network Function Virtualization
NOIT	Network Operations & Information Technology; this term network broadly includes carriers, service providers, enterprise networks, and designated/contacted 3rd party administrators.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0086

OS	Operating System
OTT	Over the Top
P-GW	Packet Gateway
PaaS	Platform-as-a-Service
PBX	Private Branch exchange (telephony)
PDG	Packet Data Gateway
PDN	Packet Data Network
PHY	Physical Layer
PLMN	Public Land Mobile Network
PnP	Plug and Play
PP	Packet Processor
PR	Platform Resources
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RA	Routing Area
RAN	Radio Access Network of cellular networks
RAPI	Resources API
RFID	Radio Frequency Identifier
RNC	Radio Network Controller
S-GW	Serving Gateway
S1-AP	S1 Interface Application Protocol (an LTE protocol between MME and eNodeB)
SaaS	Software-as-a-Software
SDK	Software Development Kit
SDN	Software Defined Network
SGSN	Serving GPRS Support Node (used in core

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0087

	network of 2G/3G systems)
SIM	Subscriber Identity Module
SiP	System in a Package
SMS	Short Message Service
SoC	System on a Chip
SON	Self-Optimizing Network
SSDP	Simple Service Discovery Protocol
TA	Target Area
TNE	Test Network Environment
TTM	Time to Market
UMTS	Universal Mobile Telecommunication System
UPnP	Universal Plug and Play
URA	UMTS Routing Area
URI	Uniform Resource Identifier
UTRAN	UMTS Terrestrial Radio Access Network
VLR	Visitor Location Register
VM	Virtual Machine
VPN	Virtual Private Network
WAN	Wide Area Network
WLAN	Wireless Local Area Network
WLAN	Wireless Local Area Network
XMPP	Extensible Messaging and Presence Protocol

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0088

[0026] Database. One or more large structured sets of persistent data maintained upon a computer system organized and structured according to a software system defining rules for organization as well as responding to queries to read, write, or modify data as well as provide statistical information regarding the contained data. As used in this disclosure in describing the dSDN, a database may be either a single unified system or a distributed system wherein certain database elements are located upon different systems or servers which may be in different physical locations, acting in harmony to appear as one unified database. Where databases are described, it will be understood by one of ordinary skill in the art that (i) alternative database structures to those described may be readily employed, and (ii) other memory structures besides databases may be readily employed. Any illustrations or descriptions of any sample databases presented herein are illustrative arrangements for stored representations of information. The database formats may include relational databases, object-based models and/or distributed databases which could be used to store and manipulate the data types described herein. Likewise, object methods or behaviors of a database can be used to implement various processes, such as the described herein. In addition, the databases may, in a known manner, be stored locally or remotely from a device which accesses data in such a database.

[0027] The term "determining" and grammatical variants thereof (e.g., to determine a price, determining a value, determine an object which meets a certain criterion) is used in an extremely broad sense. The term "determining" encompasses a wide variety of actions and therefore "determining" can include calculating, computing, processing, deriving, investigating, looking up (e.g., looking up in a table, a database or another data structure), ascertaining and the like. Also, "determining" can include receiving (e.g., receiving information), accessing (e.g., accessing data in a memory) and the like. Also, "determining" can include resolving, selecting, choosing, establishing, and the like. The term "determining" does not imply certainty or absolute precision, and therefore "determining" can include estimating, extrapolating, predicting, guessing and the like. The term "determining" does not imply that mathematical processing must be performed, and does not imply that numerical methods must be used, and does not imply that an algorithm or process is used.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0089

[0028] The functionality and/or the features of a single device that is described may be alternatively embodied by one or more other devices which are described but are not explicitly described as having such functionality/features. Thus, other embodiments need not include the described device itself, but rather can include the one or more other devices which would, in those other embodiments, have such functionality/features.

[0029] Devices that are described as in "communication" with each other or "coupled" to each other need not be in continuous communication with each other or in direct physical contact, unless expressly specified otherwise. On the contrary, such devices need only transmit to each other as necessary or desirable, and may actually refrain from exchanging data most of the time. For example, a machine in communication with or coupled with another machine via the Internet may not transmit data to the other machine for long period of time (e.g. weeks at a time). In addition, devices that are in communication with or coupled with each other may communicate directly or indirectly through one or more intermediaries.

[0030] Although process (or method) steps may be described or claimed in a particular sequential order, such processes may be configured to work in different orders. Further, some steps may be performed simultaneously despite being described or implied as occurring non-simultaneously (e.g., because one step is described after the other step) unless specifically indicated. Moreover, the illustration of a process by its depiction in a drawing does not imply that the illustrated process is exclusive of other variations and modifications thereto, does not imply that the illustrated process or any of its steps are necessary to the embodiment(s), and does not imply that the illustrated process is preferred. Where a process is described in an embodiment the process may operate without any user intervention.

[0031] It will be readily apparent to one of ordinary skill in the art that the various processes of the dSDN described herein may be implemented by, e.g., appropriately programmed general purpose computer(s), special purpose computer(s) and computing device(s). Typically a processor (e.g., one or more microprocessors, one or more microcontrollers, one or more digital signal processors) will receive instructions (e.g., from a memory or like device), and execute those instructions,

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0090

thereby performing one or more processes defined by those instructions. Instructions may be embodied in, e.g., one or more computer programs, one or more scripts.

[0032] A "processor" means one or more microprocessors, central processing units (CPUs), computing devices, controllers, microcontrollers, digital signal processors, or like devices or any combination thereof, regardless of the architecture (e.g., chip-level multiprocessing/multi-core, Reduced Instruction Set Computer (RISC), Complex Instruction Set Computer (CISC), Microprocessor without Interlocked Pipeline Stages, pipelining configuration, or simultaneous multithreading).

[0033] Further, programs that implement methods described herein may be stored and transmitted using a variety of media (e.g., computer readable media) in a number of manners. In some embodiments, hard-wired circuitry or custom hardware may be used in place of, or in combination with, some or all of the software instructions that can implement the processes of various embodiments. Thus, various combinations of hardware and software may be used instead of software only to implement the embodiments.

[0034] The term "non-transitory computer readable medium" in this disclosure refers to any medium, a plurality of the same, or a combination of different media, that participate in providing data (e.g., instructions, data structures) which may be read by a computer, a processor or a like device. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media may include, for example, optical or magnetic disks and other persistent memory. Volatile media may include dynamic random access memory (DRAM), which typically constitutes the main memory. Transmission media may include coaxial cables, copper wire and fiber optics, including the wires that comprise a system bus coupled to the processor. Transmission media may include or convey acoustic waves, light waves and electromagnetic emissions, such as those generated during radio frequency (RF) and infrared (IR) data communications. Common forms of non-transitory computer-readable media in which the dSDN may be implemented include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, any other magnetic medium, a CD-ROM, DVD, any other optical medium, punch cards, paper

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0091

tape, any other physical medium with patterns of holes, a Random Access Memory (RAM), a programmable read only memory (PROM), an erasable programmable read only memory (EPROM), a flash electrically erasable programmable read only memory (FLASH-EEPROM), any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0035] Various forms of computer readable media in this disclosure may be involved in carrying data (e.g. sequences of instructions) to a processor. For example, data may be (i) delivered from RAM to a processor; (ii) carried over a wireless transmission medium; (iii) formatted and/or transmitted according to numerous formats, standards or protocols, such as Ethernet (or IEEE 802.3), Bluetooth, and Transmission Control Protocol/Internet Protocol ("TCP/IP"), Time Division Multiple Access (TDMA), Code Division Multiple Access (CDMA), and 2G/3G; and/or (iv) encrypted to ensure privacy or prevent fraud in any of a variety of ways well known in the art.

[0036] Thus in some embodiments in this disclosure a description of a process may be a description of a non-transitory computer-readable medium storing a program for performing the process. The computer readable medium may store (in any appropriate format) those program elements which are appropriate to perform the method.

[0037] In an embodiment, a server computer, network element or centralized authority may not be necessary or desirable. For example, an embodiment may be practiced on one or more devices without a central authority. In such an embodiment, any functions described herein as performed by the server computer or data described as stored on the server computer may instead be performed by or stored on one or more such devices.

[0038] In interpreting the present application (which includes the claims), one of ordinary skill in the art shall refer to the prosecution history of the present application, but not to the prosecution history of any other patent or patent application, regardless of whether there are other patent applications that are considered related to the present application, and regardless of whether there are other patent applications that share a claim of priority with the present application.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0092

[0039] Numerous embodiments are described in the present application, and are presented for illustrative purposes only. The described embodiments are not, and are not intended to be, limiting in any sense. One of ordinary skill in the art will recognize that the disclosed embodiment(s) may be practiced with various modifications and alterations, such as structural, logical, software, and electrical modifications. Although particular features of the disclosed invention(s) may be described with reference to one or more particular embodiments and/or drawings, it should be understood that such features are not limited to usage in the one or more particular embodiments or drawings with reference to which they are described, unless expressly specified otherwise.

[0040] The present disclosure is not a literal description of all embodiments of the invention(s). Also, the present disclosure is not a listing of features of the invention(s) which must be present in all embodiments. A description of an embodiment with several components or features does not imply that all or even any of such components/features are required. On the contrary, a variety of optional components are described to illustrate the wide variety of possible embodiments of the present invention(s).

[0041] The following description has been presented for purposes of illustration and description and is not intended to be exhaustive or to limit the embodiments to the precise form disclosed. Many modifications and variations are possible in light of the teachings disclosed herein.

[0042] The embodiments were chosen and described to explain principles of operation and their practical applications. However, the scope of the invention is to be defined by the claims.

[0043] The major technical reason for inflexibility in network infrastructure has been related mostly to the proprietary hardware and rigid software architecture in the existing network elements (routers, switches, gateways, cellular base-stations, WiFi access points, etc.). With recent advancements in semiconductor device manufacturing, there is no longer a cost barrier to add computing power to many electronics in the market including networking infrastructure. The addition of such compute capabilities would open up new opportunities for programmable platforms in the network and ultimately creating more flexible network architecture

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0093

and business models.

[0044] The Distributed Software Defined Network (dSDN) disclosed herein is an end-to-end architecture that enables secure and flexible programmability across a network with full lifecycle management of services and applications (fxApp). The dSDN also harmonizes fxApp deployment across the network independent of the hardware vendor. As a result, the dSDN simplifies the network deployment lifecycle from concept to design to implementation to decommissioning. In this disclosure, the dSDN is applied to the wireless networks as an exemplary embodiment. However, the dSDN is not limited to wireless networks embodiments and it could be applied to many other network types including enterprise, wireline service providers, data centers, and Internet of Things (IoT).

[0045] Figure 1 demonstrates the network architecture of a 4G system 100. (In alternative embodiments, the system and method disclosed may also be implemented in a 2G/3G network as well). In a cellular network, the base station or BS (denoted as BTS in 2G, NodeB or NB in 3G, and eNodeB or eNB in 4G systems) is responsible for termination of a wireless link (air-interface) to the mobile device (also known as the User Equipment (UE)). The BS then transports the user data to the core network and communicates control-signaling messages with the core network. In a typical cellular network, there are thousands of BS's covering nations. Each BS covers an area known as a cell. As the demand for mobile Internet expands, the mobile network operator (MNO) needs to deploy smaller cells so it can reuse the frequency more often.

[0046] A BS is made up of the following main functional blocks: 1) radio frequency (RF) front end; 2) clock; 3) baseband; 4) power manager; and 5) central processing unit (CPU). Some BS's also have a dedicated packet processor (PP) to accelerate packet processing in hardware. BS vendors may also use custom or merchant System-on-a-Chip (SoC) or Multi-Chip-Packaging (MCP) solutions to combine various functions into a single chip (e.g., PP, CPU, and baseband).

[0047] Even though there has been attempts to open up various interfaces in the BS (such as the Open BS Architecture Initiative that defines interfaces between these functional modules in the BS), there has been very little to no efforts to unify the

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0094

programmability of the BS itself. As a result, the MNOs suffer from the following difficulties. First, the core network Internet edge (CNIE), Packet Data Gateway (PDG) and the surrounding functions are becoming extremely complicated, non-scalable, and expensive. Second, each MNO needs certain customization and feature sets. Currently, they depend on their vendors for these customizations, which could cost the MNO tremendously both financially and in regards to time-to-market (TTM); hence, hindering innovation. In addition, usually these features are put into standards or in the vendor's feature set, thus eliminating the MNO differentiation against the other MNOs. Third, if the MNO is multi-vendor in their CNIE, coordinating all the vendors to implement the same features is usually a difficult and a time-consuming effort. As a result, more add-on network appliances are introduced into the CNIE which adds to the network management complexity. Fourth, as a BS becomes Internet Protocol (IP) based, there is more visibility into the user traffic types and new innovative opportunities are missed for creative traffic shaping features, backhaul bandwidth optimization, prioritization at radio edge, power management algorithms, etc. Fifth, many features implemented in CNIE present suboptimal performance. For example, the filtering enforcement at CNIE is fundamentally inefficient since the packets have to travel all the way through the expensive air-interface, the backhaul, various other core network elements and transport networks to get to the CNIE. Sixth, the Radio Access Network (RAN) deployments are a very expensive endeavor for MNOs. The current rigid BS designs limit innovation and force the MNOs to undergo major upgrades every few years.

[0048] The problems stated above apply to many different types of networks even though here they are presented in cellular networks as a focus. In the present disclosure, the dSDN exhibits a new paradigm in the software programmability of networks that would address the problems above and enable many more advanced features.

[0049] The dSDN may be located at a network element (or system) 200 which is shown in detail in Figure 2 or in alternative embodiments the functions of the dSDN may be divided among a plurality of network elements (or systems) which are similar in structure to network element 200. Each network element 200 may comprise one or

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0095

more system control logic 202 coupled with at least one or a plurality of processor(s) 204, system memory 206, a network interface 208 (including a transceiver 208a), and input/output (I/O) devices 210. The processor(s) 204 may include one or more single-core or multi-core processors. The processor(s) 204 may include any combination of general-purpose processors and dedicated processors (e.g., graphics processors, application processors, baseband processors, etc.). System control logic 202 for one embodiment may include any suitable interface controllers to provide for any suitable interface to at least one of the processor(s) 204 and/or to any suitable device or component in a packet network in communication with system control logic 202. System control logic 202 for one embodiment may include one or more memory controller(s) to provide an interface to system memory 206. System memory 206 may be used to load and store data and/or instructions, for example, for network element 200. System memory 206 for one embodiment may include any suitable volatile memory, such as suitable dynamic random access memory (DRAM), for example. System memory 206 may also include non-volatile memory including one or more tangible, non-transitory computer-readable media used to store data and/or instructions, for example, such as the embodiments described herein with regard to the dSDN. The non-volatile memory may include flash memory, for example, and/or may include any suitable non-volatile storage device(s), such as one or more hard disk drive(s) (HDD(s)), one or more compact disk (CD) drive(s), and/or one or more digital versatile disk (DVD) drive(s). The system memory 206 may include a storage resource physically part of a device on which the network element 200 is installed or it may be accessible by, but not necessarily a part of, the device. For example, the system memory 206 may be accessed over a network via the network interface 208 and/or over Input/Output (I/O) devices 210. Network interface 208 may include a transceiver 208a to provide a radio interface for network element 200 to communicate over one or more network(s) and/or with any other suitable device. Network interface 208 may include any suitable hardware and/or firmware. The network interface 208 may include a plurality of antennas to provide a multiple input, multiple output radio interface. Network interface 208 for one embodiment may include, for example, a wired network adapter, a wireless network adapter, a telephone modem, and/or a wireless modem. For one

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0096

embodiment, at least one of the processor(s) 204 may be packaged together with logic for one or more controller(s) of system control logic 202. For one embodiment, at least one of the processor(s) 204 may be integrated on the same die with logic for one or more controller(s) of system control logic 202. In various embodiments, the I/O devices 210 may include user interfaces designed to enable user interaction with the network element or system 200, peripheral component interfaces designed to enable peripheral component interaction with the network element or system 200, and/or sensors designed to determine environmental conditions and/or location information related to the network element or system 200. In various embodiments, the peripheral component interfaces may include, but are not limited to, a non-volatile memory port, a universal serial bus (USB) port, an audio jack, and a power supply interface.

dSDN Architecture

[0050] In the dSDN architecture, the network features are virtualized and may be distributed across separate network elements that cooperate together to create an advanced, programmable, and scalable network. In an alternative embodiment, the network features may be located at a single network element. A high-level overview of a dSDN system 300 is depicted in Figure 3. The dSDN system 300 consists of a flexible network device (fxDevice) 302, a flexible cloud platform (fxCloud) 304, an application management portal (fxManager) 306, and an infrastructure application market place(fxStore) 308. The fxDevice 302 is a network device that is powered by a sandboxing operating system named flexible operating system (fxOS) 302a. The sandboxing operating system ensures each application runs as a dedicated process in sure isolation from the other applications. The fxOS 302a may be built based on an existing sandboxed operating system (OS) (e.g., Android®) by extending several aspects of such an OS with routing/networking layers and supports for operating as the wireless infrastructure. The fxDevice 302 may host several independent and securely isolated applications (on top of fxOS 302a) named here fxDeviceApp 302b.

[0051] In the dSDN system, fxDeviceApp 302b may have a sister app in the backend cloud infrastructure (i.e., flexible cloud platform 304) referenced here as fxCloudApp 304a. The fxCloudApp 304a in the cloud is paired with its fxDeviceApp 302b in the fxDevice 302. The fxCloudApp 304a and the fxDeviceApp 302b

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0097

collectively form a distributed application (dApp or fxApp). In this description, when fxDeviceApp is referenced it shall mean any application that may have software components in the fxDevice, fxCloud, or both. In general, an fxApp package may include the following software components:

fxDeviceApp binary;

fxCloudApp binary;

Manifest files; and

Signatures.

[0052] It is important to note that the fxDeviceApp 302b and fxCloudApp 304a could use any protocol to communicate with each other and there is no need for standardizing this communication allowing for ultimate freedom for the developers. This allows for the system to operate in a loosely coupled autonomous fashion allowing for asynchronous communication in a distributed fashion.

[0053] The fxApps lifecycle is managed via fxManager 306. The fxManager 306 presents a user-friendly portal to the network administrator (admin) to discover, test prior to deployment, provision, and deprovision dApps. Using fxManager 306, the admin may discover new fxApps from the fxStore 308, which presents all the tested and certified vApps as well as showing the supported fxOS version, support hardware platforms, and other information such as reviews and number of commercial deployments by all NOIT customers.

fxOS

[0054] fxDevice 302 is defined as any networking equipment embedded with a special flexible operating system (fxOS) 302a. The purpose of fxOS 302a is to enable a carrier-class programmable platform. The fxOS 302a is configured to do, but is not limited to, the following. First, since fxOS 302a is designed for network equipment, it allows for efficient execution and running all the time. A plurality of power save modes may be implemented to ensure energy saving schemes without hindering the functionality of the fxDevice 302. Second, it enables simultaneous execution of several flexible applications (fxDeviceApp 302b) on the fxDevice 302. Third, it allows for application of a legacy compiled binary or byte-code of another language (for example, Java or Android®) in addition to fxDeviceApps 302b. Fourth, it allows

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0098

FastPath processing for data plane packets. The packet forwarding in this case separated from the control plane operations. The FastPath engine allows integration with specialized hardware accelerators or packet processors (PP). The packet forwarding may be implemented with software/kernel enhancements all in software. Another FastPath implementation may use both software/kernel enhancement and hardware packet processors (PP). Fifth, the fxOS 302a is capable of Deep Packet Inspection (DPI) and identifying all the flows and their related protocols, services and present this information in a format useful to application developers. The fxOS 302a can act on the data flows as the directed by the application developer. Sixth, applications are securely isolated from each other and to communicate with them is completely configurable by policies defined by the admin. Seventh, resource utilization of each application is policed such that an application may not be able to exceed its usage. For example, an application is pre-allocated the memory size, cache/storage space, and CPU cycles. Eighth, the applications may be signed with unique certificate security keys. The security keys and certificates may be allocated and/or signed by a certificate authority. The fxDevice 302 may be protected by validating signed applications to run or to install in the device. Ninth, it allows hot upgrade of the software and applications on the platform with little to no interruption to the operational aspect of the platform and its applications. Tenth, the fxOS 302a may boot by using the secure code from the network servers (cloud-booting). This would ensure tamper-proof software foundation for the fxDevice 302. Eleventh, there may be a secure and programmable firewall in the fxOS 302a. This embedded firewall protects against attacks to the fxDevice 302 and network attached to the network. The fxOS 302a firewall may be configured by the developer and it does allow replacing of attack detection algorithms (e.g., man-in-the-middle detection algorithm). The fxOS 302a may allow secure caching of user data, webpages, and media files (video and audio) in the fxDevice 302. The secure connections to the cloud (fxCloud 304) can enable sharing and control of the caching between the fxCloud 304, fxDevice 302, or the rest of the network. Twelfth, the fxOS 302a can support multi-threading and multi-core CPUs. This feature may dynamically distribute the threads to various CPU core without developers assistance.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0099

[0055] The fxOS 302a allows secure connection to the fxCloud 304. This secure connection allows communication between fxDeviceApps 302b in the fxDevice 302 and fxCloudApp 304a in the fxCloud 304. The fxOS 302a controls the access of applications to platform APIs. An application is categorized into an Access Level (AL) based on which the fxOS 302a decides that the application does not exceed its AL and access APIs that it is not authorized. The fxCloud 304 and the fxOS 302a create a virtual fabric (fxVF) for messaging between applications. The actual message can travel directly between two fxDevices or it may traverse the fxCloud 304. The messaging complexity may be abstract for the developer.

[0056] Figure 4 demonstrates an examplary view of the software layers in the fxOS 302a. Each layer of this software stack is modified and redefined to optimize for embedded networking equipment of wireless and wireline nature. Starting from the bottom of the stack, the kernel also connects to the hardware resources via the device drivers. Examples of such resources are wireless baseband SoC 402, hardware packet processor 404, Wide Area Network (WAN) uplink interface 406, cache and storage 408, and multi-core CPU 410. The OS kernel 412 is enhanced with the FastPath function 414 to accelerate packet processing (by help of hardware features such as Intel's Data Plane Development Kit (DPDK) or pure software enhancements) and other related data path functions. The enhanced kernel with FastPath 412, 414 would enable visibility into the data flows, sniffing the Non-Access Stratum (NAS) signaling (which is the signaling between the UE and the Core Network and the RAN is supposed to be transparent to), policy enforcements, regular packet forwarding functions, collection of flow statistics into Flow Information Base (FIB), and encryption/decryption functions.

[0057] Above the kernel, there are several tools 416, native daemons 418, native libraries 420 and Hardware Abstraction Layers (HAL) 422. These tools abstract out the hardware dependencies for the upper layers and programmers. The Runtime 424 is an embedded virtual machine capable of securely isolating and executing the applications. System Services 426 are a set of services always running and available to the developers (e.g., timing and messaging services).

[0058] fx-Framework Library 428 consists of several Frameworks. The new

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0100

frameworks are the library extensions introduced explicitly for the dSDN framework. Each of these frameworks adds a set of methods of (functions) and data structures for the following examples:

Wireless Manager 428a: manipulation, monitoring, and configuration of the wireless interfaces;

System Manager 428b: configuration, provisioning, and changing major settings of the fxDevice 302;

Flow Controller 428c: allows for manipulation of data path, tunneling, switching, and optionally coordinate this activity with fxCloud 304;

Energy Management 428d: enables fine-grained control over various elements in the fxDevice 302 that could be moved into various power stares, for example: ON, OFF, SLEEP, SUSPEND;

Big Data Manager 428e: manipulation, processing, and organization of data collected from various elements. This framework could enable close coordination between the fxDevice 302 and the fxCloud 304;

Content Manager 428f: for data sharing (database sharing) between the applications within the fxDevice 302 or the fxCloud 304 and between applications in fxDevice 302 and fxCloud 304. The Content Manager 428f may use URI to address the contents to be shared. The Content Manager 428f allows access to the application data according to the policy settings and permissions;

Policy Manager 428g: responsible for execution and implementation of polices set up by the admin. The application developers could use this framework to perform queries about the permissions, limitations, and rules;

Security Manager 428h: allows for access to the security protocol libraries and rewriting some of security algorithms such as the man-in-the-middle detection algorithm;

General Framework: general compute and logic building that may be

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0101

inherited from existing general OS frameworks.

[0059] Using the frameworks, there are potentially at least two types of application types that are possible. First, the System Native Applications 430 that are provided as the initial application load into the platform. These applications could be used by other applications by links or API. Second, the fxDeviceApps 302b that are written by third party developers or the NOIT.

[0060] Figure 5 further details the FastPath 414 architecture. In this approach, the FastPath 414 is divided into the Virtual Switch (data-plane) 502 and the Switch Controller (control-plane) 504. Such separation makes possible flexible deployment wherein the function of the Switch Controller 504 could move to the fxCloud 304 by an admin decision or by the dSDN system intelligence for example in the case of failure, redundancy, stateful reboot, or Hot Upgrade. The Virtual Switch 502 may include several functions and may interact with the Switch Controller 504 in the following exemplary ways:

Firewall 502a: implements the policies given to it by the firewall logic 504a in the fxOS Daemons.

Flow/NAS Discover 502b: continuously monitors data packets and discovers new flows and NAS signaling. It would then report it to the DPI Logic 504b where the flow intelligence resides and matching to meta-data happens. DPI Logic 504b in turn updates the Flow Information Base Information (FIB) 502c of the Virtual Switch 502.

FastPath & DPI Processing 502d: forwards, routes packets, and performs DPI functions as instructed by the FIB 502c. It also collects statistics of packets and reports them to the Monitoring Logic 504c.

Policy Enforcement 502e: this function enforces various policies on the traffic such as QoS marking according to the policies received from the Policy Logic 504d.

Shaper 502f: this function could shape (queue) the traffic according to the available bandwidth on the uplink or WAN interface.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0102

Virtual Fabric (fxVF) 504e: enables transparent switching of application and system messages of a dApp between the fxCloud 304 and the fxDevice 302 and between different dApps. The fxVF 504e is an abstraction layer that hides most of the complexities of messaging from the developers. The fxVF 504e uses the policies and firewall rules to enforce security and policies.

[0061] The fxOS 302a is mostly used for networking applications. There are some potential use cases where there is a need for a general purpose OS to run the legacy software applications with no need of recompilation. An example of a double OS situation may be where a digital signage (using a general purpose OS such as Android®) in a shopping mall also acts as a smart small-cell (using the fxOS 302a). In such a solution, any existing Android® application (for example) could be loaded in the fxDevice 302 and in parallel networking applications (fxDeviceApps 302b) could run on the same fxDevice 302.

[0062] Figure 6 presents a potential software layer of the fxOS 302a co-existing alongside a General Purpose OS 602. Another implementation may be a standalone fxOS without the company of another general OS. It is important to note that the shown co-existence allows reuse of the kernel so there would be no need to re-port the drivers of hardware resources twice.

fxCloud

[0063] The fxCloud 304 is an integral part of the dSDN framework. The fxCloud 304 connects to various parts of the fxDevice 302 to form a single virtual view of the system. Examples of the main fxCloud 304 responsibilities include the following. First, it interacts and manages all applications, firmware, and fxOS 302a in the fxDevice 302. Second, it secures isolation environments for distributed apps. Third, it directs the Virtual Machine lifecycle management including load-linked VM creations and destruction (Cloud Breathing). Fourth, the fxCloud 304 executes instructions defined by the fxManager 306. Fifth, it implements uniform policy (including security) distribution and execution across dSDN 300 (including the fxDevice 302, fxDeviceApps 302b, fxCloudApps 304a, anddApps). Sixth, Plug-and-Play commissioning of thefxDevice 302 introduction

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0103

and network expansion. Seventh, Distributed Resource Service (dRS) by which distributed applications can seamlessly communicate and message to each other using published APIs or proprietary methods within the pre-defined policies. Eighth, security violation discovery and alert system. Ninth, translation and/or exposure of Resources APIs (RAPI) to the systems outside the dSDN 300 within pre-defined policies. Tenth, fxCloud 304 in coordination with the fxDevice 302 creates a Distributed Notification Service (dNS) across the dSDN 300 by which the applications could be notified of an event to wakeup and/or respond to the event. Eleventh, fxCloud 304 in coordination with the fxDevice 302 creates a Distributed Content Management Service (dCMP) which provides a virtualized and distributed database system for applications to share application-specific or system-specific information with the pre-defined policies. Twelfth, the Switch Controller 504 function in the fxCloud 304 could take control of the controller function of the fxOS 302a in the fxDevice 302. This would allow controller switching from local to remote and back. Examples of such a configuration is a temporary reason where the fxOS Switch Controller 504 may not be available during reboot or software upgrade.

[0064] In the discussion below, these services and features are expanded upon.

fxDeviceApp-0 & fxCloudApp-0

[0065] The fxDevice 302 and fxCloud 304 have embedded System Applications categorized as fxDeviceApp-0 and fxCloudApp-0, respectively. The fxDeviceApp-0 is responsible for all default and core functional capabilities of the fxDevice 302 as listed below:

Firmware of the fxDevice;

Default communication protocols (e.g. LTE S1-AP, baseband software, security protocols, GTP, DNS, tunneling protocols);

Message communication protocol between the fxDevice and fxCloud;

Initial fxDevice settings;

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0104

Boot Loader;

fxDevice X.509 digital certificate;

OS and Kernel upgrades;

Upgrade and commissioning agents for the fxDeviceApp-0 upgrades and related procedures;

Establishing and initiating the message communication protocol between the fxDevice and fxCloud;

General fxDeviceApp install and upgrades;

Access to fxDevice hardware and software information.

The fxCloudApp-0 is responsible for all default and core functional capabilities of the fxCloud 304 such as:

Upgrade and commissioning agents for the fxCloudApp-0 upgrades and related procedures;

Hypervisor upgrades;

Virtual Machine lifecycle management (e.g. creation, suspension, destruction);

fxCloud X.509 digital certificate management;

OS and kernel upgrades;

General fxCloudApp install and upgrades;

Default communication protocols (e.g. security protocols, DNS, tunneling protocols).

Each software component can be managed and updated individually. To simplify the management of these software components, each software components can be sub-numbered such as fxDeviceApp-0.1, fxDeviceApp-0.2, fxCloudApp-0.1, fxCloudApp-0.2, etc.

fxManager

[0066] The network administrator may use the fxManager 306 to control much of the dSDN system 300. The fxManager portal could use HTML5, for example, as

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0105

a frontend technology. The major functions of the fxManager 306 may include the following.

fxDevice Management:

Commissioning (plug-and-play), de-commissioning, firmware/OS upgrading/downgrading;

Suspend/resume functions;

Unification of zones and partitions for example in case of network merger;

Detailed management of fxDevices (remote secure login) where the administrator may zoom into a particular or a group of fxDevices to monitor and set management alarms;

Security policy setting.

Application lifecycle management: fxManager 306 may work, coordinate, and/or perform these tasks in conjunctions with fxCloud 304 for:

Provisioning, de-provisioning, upgrade/downgrade applications. Some applications may only have a fxDevice component (Device-Centric Apps), may only have fxCloud component (Cloud-Centric Apps), or may have both fxDevice and fxCloud components (Hybrid Apps);

Creating a Test Network Environment (TNE) where the administrator can test certain configuration and applications without impacting the production networks;

Uploading applications developed by the NOIT or their partners into the Application Repository (AR). The AR is placeholder where the application packages are downloaded and stored for testing and verification prior to installation into the live dSDN system. An application could have several components (e.g. fxDeviceApp binaries, fxCloudApp signatures, manifest files, signatures) bundled together to form the Distributed Application Package (dAP);

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0106

Browsing through the certified applications in the fxStore and downloading them to the AR;

Application license accounting services by which a secure and accurate accounting of the installed applications are performed regularly. This inventory management function allows the administrator and the 3rd party ISV to account for live and active instants of the software in the network;

Application review tool where the network admin could use to review feedbacks and comments of other NOIT (maybe anonymous reviews) and could input his/her reviews for others (maybe anonymous reviews).

fxCloud Management:

Zoning and partitioning the network for administrative purposes. Zones and partitions could optionally put into different Virtual Machine (VM);

VM load management and Cloud Breathing policies;

The configuration conflict resolutions. Some configurations applied to the resources may conflict with each other (e.g., one firewall rule may state drop Netflix packets explicitly and another rule may state allowing Netflix explicitly). The fxCloud resource manager could monitor these configurations and discover the conflicts and signal those to the network administrator.

Security Management:

Security policy setting;

Inter-App communication setting;

Security of communication between fxDevice and fxCloud;

API Access Level (AL) definitions and (re-)assignment to the applications (permissions);

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0107

Risk Analysis of the applications;

Certificate and signature verification of the applications;

Content sharing between applications and associated policies.

[0067] An fxManager Portal per the above description may have an "Account" button which can be used to create admin users and their privileges. The fxManager Portal could also include a geographical map illustrating locations of fxDevices and their status. The fxManager 306 functions could be integrated into other network management tools. The fxManager 306 could potentially build the inventory of the network devices (fxDevices) showing all the available devices and statistics. As described above, the network administrator could optionally procure the applications a third party hosted application store (fxStore). The fxStore 308 could potentially categorize the applications several groups (as shown in Figure 7 which is an example view of the fxStore portal). The fxStore 308 could also show the reviews of the advertised applications. In order to maintain quality, the third party fxStore 308 may perform rigorous testing of the applications and accompany those results with the advertised application.

fxSDK

[0068] The flexible Software Development Kit (fxSDK) may include a development environment and other tools to facilitate development, testing, debugging and verifications of applications for dSDN environments. The fxSDK is unique in several aspects exemplified as follows. First, since the dSDN applications may have a device component (fxDevice 302) and cloud component (fxCloud 304), the fxSDK allows the developer to develop the components together to simplify the development and testing. Second, the dRMS offers various system level APIs with remote accessibility. In other words, the APIs in the fxDevice 302 could be accessed via the applications in the fxCloud (fxCloudApps) and vice versa. The fxSDK simplifies the usage system created APIs and application created APIs. Third, it is essential for the dSDN to ensure quality and security of applications in the network. The fxSDK may be accompanied with an extensive Risk Analysis Utility (RAU) which verifies all the APIs that are being used by the application and based on its potential danger to the

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0108

network, it would show the risk analysis and may make specific suggestions to reduce the risk. Once completed, the risk analysis result is included in the dAP for posting to the fxStore and consumed by the fxManager user (the admin). Fourth, the fxSDK would also ensure that the applications do not exceed their planned memory and resource usage. Fifth, the Integrated Development Environment (IDE) may provide various tools to actively demonstrate errors, defined APIs, security dangers, and excessive resource usage. Sixth, the IDE could also provide a flexible Simulator (fxSimulator) for the developer test the application in a network setup in presence of other network elements.

Deployment Examples

[0069] The dSDN 300 offers flexible deployment options. In this section, a few exemplary embodiments are presented.

[0070] Figure 8 demonstrates a shared cloud deployment where several NOITs (802, 804, 806) share an fxCloud 808. In order to ensure complete isolations, each NOIT may be assigned a separate Virtual Machine (VM) 802a, 804a, 806a. A good example of this deployment option may be when multiple service providers outsource their fxCloud deployment to a third party vendor that offers a SaaS/PaaS solution. It is important to note that in this option, each NOIT should have its own secure instance of the whole system. This option could potentially create a full outsourcing model for network operations including possibilities of seamless merging of networks of two or more NOITs.

[0071] Another deployment option may be a Zoned Deployment as shown in Figure 9. In this option, a large NOIT may partition their network into a plurality of zones (Zone #1, Zone#2,… Zone#n) for administrative, software compatibility grouping, or other reason. For example, it is possible that in a large network deployment, there would be different versions of fxOS with varying capabilities in the network (the older hardware may not support newer fxOS version). In this case, the service provider may decide to group the fxDevices per fxOS versions for management simplifications.

[0072] Figure 10 is another embodiment deployed in a Programmable & Virtualized Cellular Network. In this example, the dSDN is applied to the cellular

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0109

networks (e.g., LTE). As shown, the fxDevice 1002 here is the LTE eNB (4G base-station) and the fxCloud 1004 hosts the LTE core network (EPC) functions such as MME, S-GW, and P-GW in separate VMs. Within each EPC function, new features could be added using the dSDN framework. For example, a new application to optimize signaling for Machine-2-Machine (M2M) devices could be loaded into the fxDevice 1002 and the MME function in the fxCloud 1004.

PRINCIPAL PROCEDURES & SERVICES

[0073] The dSDN framework could offer extensive services and capabilities that would simplify the programmability of the network. In this section below, a set of main procedures and services are presented as further exemplary embodiments.

fxDevice Commissioning Provisioning

[0074] Figure 11 illustrates a signaling flow chart for a fxDevice Plug-and-Play Commissioning Procedure. This procedure enables fxDevice commissioning and initial setup to provide a seamless and Plug-and-Play (PnP) deployment. The fxDevice relies on the fxDeviceApp-0 and the fxCloudApp-0 to connect to the backend management system and perform the commissioning procedure. In step 1101, the fxDevice 1120 tries to discover the fxManager 1130 using the fxManager Fully Qualified Domain Name (FQDN). The fxDevice 1120 is configured to always look for the same FQDN. If the NOIT has a private cloud deployment, the DNS server in the NOIT will be configured to point to its own fxManager IP address. If the cloud is shared, the NOIT will configure their DNS server to point to the shared cloud IP address. In step 1102, the fxDevice 1120 receives the IP address of the fxManager 1130. In step 1103, the mutual authentication may be performed where the fxDevice 1120 authenticates the fxManager 1130 and vice versa. In step 1104, a secure tunnel may be created between the fxDevice 1120 and the fxManager 1130. In step 1105, the fxDevice 1120 may optionally request for the latest firmware/OS version. In this phase, the fxDevice 1120 informs the fxManager 1130 of its hardware information (e.g. vendor, model number, CPU model/speed, memory/cache size) and software information (e.g., current firmware version, current OS version, fxDeviceApp-0 version package stating versions of fxDeviceApp-0 components). In step 1106, the fxManager 1120 may send the latest firmware/OS version (or point

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0110

the fxDevice to the right URL to download the latest load). In step 1107, if a new firmware/OS load is downloaded and verified, the fxDevice 1120 installs it. In step 1108, the fxDevice 1120 requests the default settings and the default application package (the initial applications that the NOIT admin would like to install in fxDevice 1120 prior to other applications). In step 1109, the fxManager 1130 prepares the fxCloud 1140 for the new fxDevice 1120 integration into the network. In step 1110, once completed, the fxCloud 1140 informs the fxManager 1130 of its readiness. In step 1111, the fxManager 1120 sends all the default settings and the default application package to the fxDevice 1120. In step 1112, the fxDevice 1120 then sets up a secure connection to the fxCloud 1140 (more precisely the fxCloudApp-0). In step 1113, a commissioning completion message is set from the fxCloud 1140 to the fxManager 1130. In step 1114, the fxManager 1130 signal to the fxDevice 1120 the completion of the commissioning and starts th regular operation of the fxDevice 1120. At this phase, fxManager 1130 can optionally modify the settings in the security connections between fxDeviceApp-0 and fx-CloudApp-0 and between fxDeviceApp-0 and fxManager 1130.

Application Provisioning Procedure

[0075] The Application Provisioning Procedure enables provisioning and de-provisioning of applications across the dSDN system. For simplicity, the term provisioning in this disclosure is used to present all similar procedures of provisioning and de-provisioning. The application provisioning procedures are usually triggered by the network administrator and are orchestrated by the fxManager, which works in collaboration with the fxDeviceApp-0 and fxCloudApp-0.

[0076] Figure 12 shows an example procedure for Application (dApp) Provisioning. In step 1201, the network admin identifies the applications that would need to be installed in the dSDN system. These applications are already downloaded (from the fxStore or other sources) into the Application Repository (AR). Prior to installation in the live network, the authorized network admin 1220 has probably performed rigorous testing of the application in the test network. In step 1202, the fxManager 1230 performs a detailed analysis on the potential risks that this application may

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0111

create on various parts of the network and other applications. The analysis tool uses the secure manifest file and other information to discover, for example, all the APIs used by the application, the required access-level authorization level for each of those APIs, requirements on the fxVF (inter-app, intra-app communication needs), and usage of platform resources including conflict discovery and resolutions. In step 1203, the admin defines the deployment scope that the application would be applied. Examples of deployment scopes could be a given zip code, a city, an administrative domain (zone), a shopping mall, an office building, or manual handpicking of sites or fxDevices 1240. The fxManager 1230 verifies compatibility of the OS version and resouces requirements of the application and the fxDevices 1240. If the fxManager 1230 discovers incompatibility, it informs the admin and request adjustment to the deployment scope or proposes upgrading incompatible fxDevice if possible. In step 1204, the admin 1220 initiates the installation process. In step 1205, the fxManager 1230 unpacks the dApp and if there is fxDeviceApp component, it would prepare it for submission to the fxDeviceApp-0 of fxDevices in the deployment scope. In step 1206, installation commands are sent to the fxDeviceApp-0s in a secure connection already set up as part of the fxDevice Commissioning. In step 1207, the fxDevice 1240 verifies the fxDeviceApp package and performs necessary integrity checks of the received software package. Once verified, the fxDeviceApp is installed and re-verified. In step 1208, the completion of successful install is sent to the fxManager 1230 by the fxDevice 1240. In step 1209, the fxManager 1230 prepares the fxCloudApp component of the dApp (if any). In step 1210, installation commands are sent to the fxCloudApp-0s in a existing secure connection. In step 1211, the fxCloud 1250 verifies the fxCloudApp package and performs necessary integrity checks of the received software package. Once verified, the fxCloudApp is installed and re-verified. In step 1212, the completion of successful install is sent to the fxManager 1230 by the fxCloud 1250. In step 1213, the fxManager 1230 informs the admin 1220 of the successful installation process.

Hot Upgrade Procedure

[0077] The Hot Upgrade refers to a procedure by which the software upgrades

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0112

on the system have no or minimal implication on the functionalities offered by the software component subject to the upgrade process. Herein are presented two types of software upgrades. For simplicity, the use of the term upgraded in this disclosure represents both upgrade and downgrade since they both use identical processes. First, upgrade of applications on the fxDevice (i.e., fxDeviceApp) and fxCloud (i.e., fxCloudApp). Second, upgrade of core functions in the fxDevice (aka fxDeviceApp-0) or fxCloud (aka fxCloudApp-0).

[0078] The upgrade procedures for fxDevice and fxCloud are similar. Here, the focus is on the fxDevice upgrade process since it is technically more challenging due to stricter resource limitations in the fxDevice.

General Application Upgrade

[0079] The applications could be upgraded automatically or manually by the network admin. In either case, the installation of the upgrade follows a similar procedure as the Application Provisioning Procedure. Figure 13 presents the state of the machine for the application upgrade process. In state 1301, the old version of the application is still running. In Load state 1302: the new version of the application is downloaded in the local storage while the old version is still running. In Install state 1303: the new version is installed (in the memory) while the old version is still running. In Run state 1304: the new version is activated taking charge of all the related data and application state and the old version is deactivated while still remaining in the local storage. In Commit state 1305: the old version is fully removed after ensuring the new version is running properly as expected.

fxDeviceApp-0 Upgrade

[0080] The fxDeviceApp-0 is a collection of software components that perform the core functionalities of the fxDevice. As a result, this hot upgrade is quite challenging. The most important thing is to ensure that major functionalities of the fxDevice (such as packet forwarding) remain

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0113

intact during the upgrade process. The fxDeviceApp-0 upgrades may be categorized as:

- No OS Kernel upgrade; or
- OS Kernel upgrade.

No Kernel Upgrade

[0081] As described earlier, the packet forwarding function could be broken into the Virtual Switch and the Switch Controller. In order to maintain, the packet forwarding function, the Switch Controller function could be performed by another device or a second processor in the same device while the fxDevice is being upgraded. Figures 14A-14C demonstrate a few examples of distributed Switch Controller function during the upgrade. In some cases, it may be possible for the fxCloud to take over the Switch Controller function (Figure 14A), or a neighboring fxDevice (Figure 14B), or second CPU in the same fxDevice (Figure 14C). The actual software upgrade procedure would be similar to Application Provisioning Procedure and the state machines are the same as the one shown in Figure 13.

Kernel Upgrade

[0082] In this case, the packet forwarding process is mostly unavailable. Therefore, one potential solution might be for the fxDevice to redirect the entire bit stream to a neighboring fxDevice. In other words, the fxDevice would have just a bare minimum routing function working and the rest of the function would be performed by a neighboring fxDevice.

[0083] All the other upgrade procedures and state machines are similar to the ones performed for the No Kernel Upgrade process.

Virtual Fabric Service

[0084] The Virtual Fabric (fxVF) provides an abstraction layer for application to communicate with each other whether they are in the fxDevice or fxCloud. Various frameworks and services may use the fxVF service. Figure 15 illustrates fxVF 1502a, 1503a, 1504a in the fxDevices 1502 and 1503 and fxCloud 1504. The fxVF 1502a,

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0114

1503a, 1504a may use Extensible Messaging and Presence Protocol (XMPP) for this messaging as the default protocol. It is important that the developers may use their own communication protocols between the fxDeviceApps 1502b, 1503b and the associated fxCloudApp 1504b. fxVF 1502a, 1503a, 1504a provide a secure routing mechanism.

Intra-App Messaging

[0085] This is an example of where one application in the fxDevice communicates with the same instance of the app in another instance of the same application in another fxDevice. The actual messages could go directly between the fxDevice or via the fxCloud. As an example, in the case of mobile networks, this messaging could be used to transfer user specific context from one eNB to another as a user hands off to a new eNB. In LTE, Private Messages on X2 interface could be used for messaging between the eNBs that act as the fxDevices.

Inter-App Messaging

[0086] In some cases, different applications may need to communicate with each other via their published APIs. In this case, the security policies set up by the network administrator determines which applications could communicate with each other for what purpose. The fxVF follows the security policies determined by the network administrator for inter-application communications.

distributed Resources Service (dRS)

[0087] The resources could be platform resources or APIs offered by the applications. The fxCloud and fxDevice resources could include (as examples):

Firewall;

Data and statistics;

Storage (usually abundant at fxCloud);

Compute (usually abundant at fxCloud);

Load evaluator;

General settings;

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0115

Routing engine (usually relevant to the fxDevice);

Virtual Machine master controller (usually relevant to the fxCloud);

Power controller (usually relevant to the fxDevice); and

Wireless engine (usually relevant to the fxDevice).

[0088] An app (fxDeviceApp or fxCloudApp) could expose APIs to be used by the other apps. The inter-app communication is enabled by the fxVF where the policy and security provisions are enforced. The APIs exposed may be RESTful (representational state transfer) and could travel across physical network elements. Since the fxManager defines fxVF policies, the administrator could ultimately specify which APIs between which apps could communicate with each other. Figure 16 demonstrates the logical interfaces between the platform resources and applications. As shown in Figure 16, the distributed Resources Service (dRS) 1602, 1603 comprises software agents residing in the fxDevice 1604 and the fxCloud 1606. These dRS agents manage the access to the platform resolutions (including potential configuration conflict resolution) as well as facilitating inter-app APIs. The dRS 1602, 1603 allows or disallows access to platform resource or inter-app communication according to the policies defined by the network administrator using the fxManager. In summary, the dRS 1602, 1603 provides services to application developers for:

Exposing APIs to other applications;

Configuring and managing platform resource;

Policy enforcement and authorization of applications access to platform resource and other app's APIs; and

Policy conflict resolution.

distributed Content Service (dCS)

[0089] The distributed Content Service (dCS) allows the developers to seamlessly store and share the contents generated by one application with other applications and its associated application in the cloud. The dCS simplifies access to the data and brings in storage virtualization to the applications. In other words, the developers no longer would need to know where the data is actually stored (in

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0116

the cloud or on the device) and would be able to access them easily. The dCS implementation may use Virtual Fabric (fxVF) and distributed Resources Service (dRS).

distributed Notification Service (dNS)

[0090] The distributed Notification Service (dNS) is another potential tool for developers that could wake or ping an application when a particular event has occurred. For example, a load monitoring application could be notified when the CPU load on a particular fxDevice (or a target area) exceeds a certain threshold. In turn, such an exemplary application could make smart decisions on reducing the load on the CPU by forcing handoffs of users to neighboring cells (i.e. fxDevices). The dCS implementation may use Virtual Fabric (fxVF), distributed Resources Service (dRS) and distributed Content Service (dCS).

Cloud Breathing Procedures

[0091] In general, resources (compute, storage) at the cloud are more abundant. However, software licensing costs and other limitation may require smart management of resources. The Cloud Breathing, here, is defined as a mechanism where the cloud resources are automatically expanded as load increases on the system or reduced as the load decreases. This creates an automatic elasticity in the cloud dimensions. Figures 17A and 17B demonstrate procedures for the cloud breathing. Figure 17A is expansion and Figure 17B is reduction. The Load Controller in the dRS 1602, 1603 could be used to monitor the load (step 1702) on the existing VMs (e.g. CPU or memory utilization loads). The VM Master of the fxCloudApp-0 1608 could effectively act on the decisions made by the dRS's Load Controller. Again, the network administrator via the fxManager defines the policies and thresholds for such decisions.

Security Aspects

[0092] The security aspects of this solution are of outmost importance in order to ensure quality in the product networks.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0117

Identity

[0093] To create a security model, the dSDN network elements should have unique identities. The following identities may be defined:

fxDevice Identity (DID): this identity is unique globally and may be allocated at manufacturing and included in the X.509 digital certificate of the fxDevice;

fxCloud Identity (CID): this identity is unique to the NOIT domain and may be allocated by the network administrator and included in the X.509 digital certificate of the fxCloud;

fxDeviceApp Identity (DAID): this identity may be globally unique and may be allocated by the fxSDK at the point of code creation by a globally accessible server;

fxCloudApp Identity (CAID): this identity may be globally unique and may be allocated by the fxSDK at the point of code creation by a globally accessible server;

dApp Identity (AID): This identity may be globally unique and may be constructed by concatenation of DAID and CAID together;

Application Developer Identity (ADID): this is a globally unique identity allocated by a globally accessible server and may be included in the developer's X.509 digital certificate and is used to sign the final developed applications.

[0094] For addressability purposes, the identities may be presented in a URI format using the FQDN of the NOIT.

Access Level

[0095] The applications are given different access levels. The AL is used by the dRS to determine which APIs or class of APIs an application can access (e.g., a fxDeviceApp running on fxDevice should not be able normally to reboot the fxDevice). The required AL is generated by the fxSDK and included in the manifest files. It may also be published in the fxStore for that give application. A preset of ALs could be defined to categorize the applications security risks.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0118

Software Security Verification

[0096] It is important for the target network element to verify the authenticity and integrity loaded software prior to install. The uniquely defined security keys of the application developers sign the applications. The following rules may be followed to ensure software security verification:

Only the verified OS/Firmware software can install on verified hardware;

Only the verified OS/Firmware software can run on verified hardware (continuous verification of essential parts of the software);

Only the verified fxDeviceApp software can install on verified hardware;

Only the verified fxDeviceApp software can run on verified hardware (continuous verification of essential parts of the software).

[0097] Each fxDeviceApp is accompanied by a manifest that is generated by the Integrated Device Electronics (IDE) at the time of compilation of the fxDeviceApp. The manifest file specifies which libraries and frameworks the fxDeviceApp uses. The fxManager uses this information to determine the security risk of the fxDeviceApp as well as the required Access Level (AL) to run this fxDeviceApp. In order to ensure that manifest file is actually genuine manifestation of the fxDeviceApp and its integrity remains intact, a hashing algorithm could be used to generate a signature. An example is shown below:

Signature = HMAC-SHA256(K, Message);

where

Message = fxDeviceApp + Manifest

K = shared secret key between fxStore and fxManager.

[0098] The signature is generated by the IDE/fxSDK at the compile/build phase and is packaged with application binary and the manifest file. Once the distributed application package (dAP) is downloaded into the fxManager (Application Repository), the components of the signature, application binary, and manifest file are unpackaged. The fxManager then uses the same hashing algorithm to calculate

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0119

the signature as above using the pre-shared key (K). If the calculated signature and the unpackaged signature match, it would prove the authenticity of the manifest file and the integrity of both manifest file and the fxDeviceApp.

[0099] Code Generation Process (Compilation):

fxDeviceApp compiled → Manifest file generation → Signature calculation → Create fxDeviceApp Package

[0100] Install Process (Consumption):

Un-package fxDeviceApp → Signature calculation → Signature comparison

[0101] This model works if the key K can be shared with two trusted parties (fxManager and fxStore). The manifest file that is generated by the fxSDK shall be secured by integrity checks to ensure it remains intact throughout the transaction (i.e. from compilation to consumption).

Communication Security

[0102] Security mechanisms built into the fxDeviceApp-0 and fxCloudApp-0 forces the fxDevice to only use the configured fxCloud servers for Software Security Verification, fxDeviceApp downloads, and fxDeviceApp-fxCloud communication. Optionally, the fxDeviceApp-fxCloud communication can be customized per application requirements. One implementation of this interface could use SPDY Protocol or per-application VPN.

[0103] For the purpose of inter-App communication, an Inter-App Communication (IAC) policy template is applied to an fxDeviceApp that defied the communication policies between the fxDeviceApps and fxCloudApps. The network administrator may use one of the pre-configured IAC policy templates (exemplified below):

Isolated: this type of applications cannot communicate with any other application in the same fxDevice and does not have any associated fxCloudApp;

Isolated-Extended: this type of fxDeviceApps cannot communicate with any other fxDeviceApps in the same fxDevice but could communicate with its associated fxCloudApp;

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0120

Private: a private fxDeviceApp in one fxDevice could communicate with same private fxDeviceApps in other fxDevice but does not have any associated fxCloudApp;

Private-Extended: a Private fxDeviceApps in one fxDevice could communicate with same private fxDeviceApp in other fxDevice and has its associated fxCloudApp;

Community: a group of fxDeviceApps that belong to the same community can communicate with each other;

Promiscuous: a Promiscuous app could communicate with any other Promiscuous application and Community application;

Custom: the network administrator could create custom IAC Policy templates and could include (as examples):

Template name;

Inherited pre-configured template if any;

The protocols allowed to use for communication;

The maximum data rate allowed for inter-app communication between a pair of fxDeviceApps or between an fxDeviceApps and an fxCloudApp;

Policies related to the other fxDeviceApps in the same fxDevice, in other fxDevice, associated fxCloudApp and unassociated fxCloudApps;
Policies related to direct fxDeviceApp to fxDeviceApp communication or relayed via fxCloud communication.

Secure Boot

[0104] The secure boot refers to a capability where the main boot code is in the fxCloud with the following exemplary procedure:

The Boot Loader (BL) in the fxDeviceApp-0 of the fxDevice looks for a particular DNS (DNSSEC) name of the fxCloud. The network administrator should define the DNS entry in their network. If there is a private fxCloud, the resolved domain name points to the fxCloud. If using shared cloud, the resolved

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0121

domain name points to the cloud point of presence;

A secure tunnel (TLS) is created between the BL in fxDevice and the fxCloud using pre-burned X.509certificate;

BL identify itself to the fxCloud and include System Information;

fxCloud determines the appropriate boot file and points the BL to the exact address of the boot file in the fxCloud;

BL downloads the correct boot file;

BL runs Software Security Verification;

BL reboots the fxDevice and the system will be ready for operation.

Frameworks

[0105] The following describes an example of the API Framework for the dSDN. Some APIs are purely local to the fxDevice or the fxCloud and some may be extended from the fxDevice all the way to the fxCloud using dRS.

General Framework

[0106] This framework refers to general libraries and APIs inherited from the legacy OS. An example of the legacy OS may be Embedded Android®. This framework allows for general processing and may be used for algorithmic applications.

Wireless Framework

[0107] The Wireless Framework adds a set up functions related to managing and controlling the wireless module. Table 1 highlights some examples of methods/functions that may be available in the Wireless Framework.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0122

Table 1: Wireless Framework	
Method	Description
discoverNeighbors	Discovers wireless neighbors and returns their identities
retrieveNeighborInfo	Retrieves detail information of neighboring BS/AP information. This method may transparently call cloud servers to gather more information
measureInterference	Measures wireless interference observed on a frequency or a set of frequency channels
setWirelessConfig	Sets various parameters in the wireless module such as power, frequency channels, Multiple Input/Multiple Output (MIMO) antenna configuration, etc.
definePHYReceiverAlgorithm	Defines the physical layer receiver algorithm
connectedDevice	Returns details about the connected mobile devices (length of connection, data tx/rx) to a given wireless interface (channel, band, sector)
forceOffload	Offloads connected mobile devices to other access technologies (e.g. from cellular to WiFi)
homeAreaNetworking	Control of home automation wireless technologies such as ZigBee or Z-Wave
setNeighborCells	Sets neighboring cellsinformation
forceHandoff	Handoffs connected mobile devices to the same access technology but different channel or frequency
topMovers	Identifies the mobile devices with highest degree of mobility

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0123

positionDevice	Accurate positioning of mobile devices which may include geo-fencing data
borrowSpectrum	Allows one BS to borrow spectrum from another
configureMACScheduler	Configures the MAC scheduler (for example prioritize a particular user)
loadBalance	Enables load balancing across radio channels
defineMACScheduler	Defines the MAC scheduler algorithm

Security Framework

[0108] The Security Framework allows the developers to use a specialized security functions/methods. Some of these new methods may use the legacy OS frameworks. Table 2 highlights some examples of methods/functions that may be available in the Security Framework.

Table 2: Security Framework

Method	Description
setDOSDetect	Detects DOS attacks on user plane as requested
configureFirewall	Provides firewall functions such as traffic filtering, permissions (this may use dRS)
defineDOSDetectAlgorithm	Allows replacing and redefining the DOS algorithm
encryptTraffic	Encrypts the content of a message or a class of traffic as requested
authenticateUser	Authenticate users asrequested
integrityCheckContent	Verifies authenticity of a message using its checksum information
configuresIPS	Configures IPS functions on the user plane

FastPath Framework

[0109] The FastPath Framework allows the developers to identify and manipulate the data path with a DPI capability. Table 3 highlights some

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0124

examples of methods/functions that may be available in the FastPath framework.

Table 3: FastPath Framework	
Method	Description
matchTraffic	Identifies and matches the traffic type requested by the developer
detectCongestion	Detects if there is congestion on the data path
identifyTopURLs	Identifies the most visited URLs by the users
detectLocalConversation	Detect conversations and packet routing between the end user on the same fxDevice
detectDevice	Detects traffic of a given mobile user device
redirectTraffic	Redirect data traffic to a given network element. The traffic may be identified with a user or a class of traffic
connectedDevices	Returns all the connected mobile devices to a given fxDevice
assignAddress	Allocates a particular IP address to the mobile device
setPriority	Set traffic priority to a class of traffic
markTraffic	Marks the user traffic with the desired QoS tags
setupTunnel	Sets up data tunnel from fxDevice to a destination
Adscript	Adds a JavaScript to an active transiting web page
detectIdenticalFlows	Identifies the identical flows coming from the WAN/uplink interface on the fxDevice
aggregateFlows	Enables aggregation of flows to save WAN/uplink utilization
routeTunnel	Routes a packet or a class of traffic onto an established tunnel
identifyTopDevices	Identifies users that generate most traffic load

Messaging Framework

[0110] The Messaging Framework allows the developers to send messaging between the applications residing in the fxDevice and fxCloud. Table 4 highlights some examples of methods/functions that may be available in the Messaging Framework.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0125

Table 4: Messaging Framework	
Method	Description
sendIntraApp	Sends an Intra-App message from one fxDevice to another
sendInterApp	Sends an InterApp message from one application to another
sendfxCloud	Sends a message to the fxCloud
sendSMS	Sends an SMS to the user of choice

Caching Framework

[0111] The Caching Framework allows the developers to cache particular contents or files that are accessed frequently in the caching engine. The caching can be done locally in the fxDevice, clustered cache (shared amongst a few fxDevice), or cloud caching. Table 5 highlights some examples of methods/functions that may be available in the Caching Framework.

Table 5: Caching Framework

Method	Description
storeForward	Enables Store and forward model for a target class of traffic where the target traffic is matched and stored upon a condition and forwarded the condition is relieved
cacheWebSite	Caches the most visited web sites
cacheVideo	Caches the most viewed videos
cacheCloud	Caches the requested content or traffic type in the cloud storage in fxCloud

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0126

Management Framework

[0112] The Management Framework enables the developers to manage the BS platforms and perform administrative procedures. Table 6 highlights some examples of methods/functions that may be available in the Management Framework.

Table 6: Management Framework

Method	Description
upgradeFirmware	Started and executes the firmware upgrade procedure on the fxDevice
upgradeOS	Started and executes the operating system upgrade procedure on the fxDevice
remoteBootup	Performs remote boot up process where the fxDevice works closely with the fxCloud to perform the bootup
remotePower	Allows fxCloudApps to remotely turn on and off a fxDevice or elements in the fxDevice (e.g. the wireless)
gatherStatistics	Instructs fxDevice to gather specific statistics. Various commands may be consolidated for gathering statistics. Number of active users Stats of time spent by users in the cell Stats of quality perceived by the users Neighboring cell info Overall traffic passed through the system Total number of connected devices per frequency, per carrier, per site, per location Handover failure statistics Cloud aggregated data processing and full report Total traffic passed Applications used (meta-data correlated) Bandwidth consumed User's mobility pattern

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0127

setMonitoring	Allows the developer to get reports and statistics of the fxDevice traffics, errors, and neighboring environment. The reporting criteria could be set up as: 1) Period, 2) Event based (reports only if a threshold is passed; a threshold could be based on absolute numbers or deltas), 3) One-time (i.e. the information is pulled once). The exact parameters to get report on may use the same data structures used by the method “getStatistics”.

getSystemInfo	Retrieves static and dynamic system info such as: • Hardware info: vendor, model number, capability profile. GPS, Cellular, WAN interfaces, RFID/NFC • Heat temperature • Vendor Specific Information • Software info: OS version, firmware information • CPU load, per-application load (CPU, memory, accelerators) • Loaded apps and status • Memory utilization • Storage utilization • Interface status & utilization (RF, backhaul, management...)
getfxDeviceLocation	Retrieves location of the fxDevice • Location would be presented in GPS coordinates and civil address • Location enhancement: Cloud, map data, and neighbor discovery could help more accurate positioning
getPowerConsumption Rate	Retrieves current power consumption rate of the fxDevice
configureWAN	Configures the WAN settings
retrieveNeighborInfo	Retrieves neighbor information from the fxCloud
getRFIDInfo	Returns the RFID/NFC information of the fxDevice
setRFIDInfo	Sets the RFID/NFC information of the fxDevice

Extensible API Framework

[0113] This framework allows the developer to create custom APIs and make it accessible for other applications (in the fxDevice or in the fxCloud). In turn, the fxCloud could present these APIs using e.g. REST technologies (via the fxCloud

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0128

Northbound Interface) to the developer outside the dSND system. There might be limits put on the APIs exposed through the cloud to avoid potential misuse or security threats. The dRS enforces the security policies defined by the admin (via fxManager). The applications take a role of Client Application (CA) or Server Application Role (SA). The CA and SA could be distributed in the fxDevice or fxCloud. The CA makes requests and SA serves the requests.

[0114] The Extensible APIs could be categorized (as examples) into:

Read-Only Data: this type of APIs allows a CA to read data from SAs without the ability to change or request any particular action from the SAs

Read-Write Data: this type of APIs allows CA application to read data from SAs with the ability to change some of their data but no option for requesting any particular action from SAs;

Procedural: this type of APIs allows CAs to execute a particular procedure in SAs. The rights to change any data still depends on whether the API allows Read-Only Data or Read-Write Data types.

Cloud Management Framework

[0115] The Cloud Management Framework presents a collection of methods enables management of the cloud services and resource. Table 7 highlights some examples of methods/functions that may be available in the Cloud Management Framework.

Table 7: Cloud Management Framework

Method	Description
detectCloudCongestion	Detects congestions on a particular VM
deleteVM	Deletes an active VM permanently
suspendVM	Suspends an active VM but doesn't remove it
activateVM	Activates an already created VM
createVM	Creates a new VM but doesn't activate it

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0129

EXAMPLES OF DAPPS

[0116] The dSDN create an end-to-end programming platform and the possibilities of vApps are only limited by the developers' imagination. An example list of potential vApps are presented in Table 8 below.

Table 8 highlights examples of use cases possible by the Distributed Software Defined Network (dSDN)

Use	Description	Framework {Methods} used
1	Distributed SON: For example, the WiFi AP could search for the least crowded channel to minimize the interference.	Wireless {discoverNeighbor, retrieveNeighborInfo, measure nterference, setWirelessConfig, loadBalance}, General Framework, Messaging {sendIntraApp}
2	BS Hot Firmware Upgrade without impact to the basic functions of the BS.	Management {upgradeFirmware}
3	BS Hot infOS Upgrade without impact to the basic functions of the BS. For this function to support full hot upgrade without operational impact, the full platform should have dual processor.	Management {upgradeOS}
4	Secure network---based booting. This allows to keep the code securely in the cloud to avoid tampering by the ODM/CM or Man---in---the---Middle attacks.	Management {remoteBoot}
5	Turning on or off APs if nobody is in the office. Sensors in the building send information to fxCloudApp which in turn orders power control of the Aps.	Management {remotePower}
6	Store and forward model for low priority M2M applications where the target traffic is matched and stored in the case of network congestion and forwarded once the network congestion is relieved.	DataPath {matchTraffic, detectCongestion}, Cache {storeForward}
7	Replacing MAC Scheduler	Wireless {defineMacScheduler}
8	Replacing Receiver Algorithms	Wireless {definePHYReceiverAlgorithm}
9	Caching most visited website or video clip in the enterprise	DataPath {identifyTopURL}, Cache {cacheWebSite}
10	The most viewed videos and webcasts are cached in the wireless edge.	DataPath {identifyTopVideos}, Cache {cacheVideo}
11	Local Routing: data and voice sessions between two parties on the same cell are routed within the same cell bypassing the core Network.	DataPath {detectLocalConversation, RouteLocal}
12	Local Routing: data and voice sessions between two parties on the nearby cells are routed locally bypassing the core network.	DataPath {detectConversationDetect, routeLocal}
13	Temporary redirection of a data flow for monitoring and security reasons. In this use case, a particular user's traffic is tracked by redirection to a monitoring station. If user is mobile, such profile of tracking and associated redirection is transferred from fxDevice to fxDevice either directly or via the fxCloud.	DataPath {detectDevice, matchTraffic, redirectTraffic}
14	Power Calendar: On---demand coverage (triggered by passing traffic thresholds in certain cells or by calendaring): Operators need to design their network for peak usage rates at high costs, leaving their network underutilized most of the time.	DataPath {detectCongestion}, Management {remotePower}
15	DOS attack recognition and action could be blocking the user, the app, and/or notifying the user via an SMS message.	Security {setDOSDetect, configureFirewall}, Messaging {sendfxCloud}
16	Applying certain ACLs to the BS. For example, the students in the class can only access Facebook between the breaks and access to Facebook is blocked during the class.	Security {configureFirewall}
17	Forced offloading of some cellular users (users that have attached to the same cell for a while) to WiFi by changing "something" in the cellular connection to force WiFi connection (Heterogeneous).	Wireless {identifyConnectedDevice, forceOffload}
18	Service continuity and anchoring at eNB for WiFi using the same LTE IP address (Heterogeneous).	DataPath {identifyConnectedDevices, assignAddress}
19	Graphic/general processing for special use cases such as City BS	General Frameworks
20	Application filtering/throttling (secondary conditions: time---based, location---based, subscriber---type, none).	DataPath {matchTraffic}, Security {configureFirewall}
21	Local Services: The enterprise user can discover Bonjour or UPnP devices and services (e.g. printer) even when connected to LTE/3G small---cell in the enterprise. In this case, the small cell or fxDevice pro---actively listens to Bonjour and UPnP discovery and advertisement messages (SSDP) and cache the available services until a device request for such information. In the case of Bonjour, fxDevice could listen to Multicast DNS (mDNS) messages sent over cellular connection (LTE/3G) and multicast them on the LAN or WLAN interfaces.	DataPath {matchTraffic, routeLocal}
22	Local Services: The premise---owner (e.g. airport, hotel) can provide ads under the browser for sponsored WiFi or cellular (3G/LTE) access over multi---mode small cell. The premise owner could use this service to advertise special offers or third party ads. This technique uses HTML <script> for example to add special content to the bottom of the page.	DataPath {matchTraffic, addScript}

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0130

23	Distributed Content Distribution Network (CDN). In this use case, the fxDevices function as CDN that bring the contact closer to the end users.	DataPath {matchTraffic}, Cache {cacheVideo}
24	In case of national security, the service provider may limit the access to the network so only the law enforcement can use the cellular network. In this case, the admin defines a Target Area (TA) using an fxCloudApp in the fxCloud where the emergency lockdown needs to be applied. The IMSI/MSISDN number of law enforcement mobile devices are sent to the fxDevice in the TA which in turn enforce the policy. The policy may allow the public to just sent SMS while the law enforcement could have full access and priority to the system.	DataPath {matchTraffic, setPriority}, Security {configureFirewall}
25	Emails and web browsing monitored by the cloud services assisted by the agent in the dSDN network. In this case, the fxDevice detects certain traffic (email or web in this case) and redirects them to a cloud service for security cleansing.	DataPath {matchTraffic, redirectTraffic}
26	Analytics pre---analysis: The data collected from M2M devices could be processed (noise could be filtered out) and packaged (compressed) for consumption by cloud services.	DataPath {matchTraffic}, General Framework, Messaging {sendfxCloud}
27	Certain traffic could be marked with QoS DSCP codes for further processing in the network.	DataPath {matchTraffic, markTraffic}
28	Proprietary SP services and applications that run purely between the phone and the BS . For example, the SP may implement a certain application for the law enforcement to get direct feed of surveillance cameras in a target area (TA).	DataPath {matchTraffic}, Security {configureFirewall}, General Framework
29	IP---PBX implementation in the small---cell or WiFi AP for enterprise use. This would allow a simple out of a box solution. The VoIP phones connect to the fxDevice via the LAN connections. The WAN connection could connect to an fxCloudApp for PSTN calls.	DataPath {matchTraffic}, General Framework
30	Network management system can add its own agent/probe to collect and consolidate data and transmit it to its cloud element using any protocol.	Management {gatherStatistics, getSystemInfo, control}, General Framework
31	Application based performance monitoring on the wireless link (number of packets sent in idle mode, retransmit rate…). Such information could be shared with the app developer to optimize its application.	Management {gatherStatistics}, General Framework
32	Seamless network sharing and mergers. This is the case where two Network Operators would like to share or merge their network. Using dSDN's infManager, the network administrators could merge the network administration control. This process may include setting up the desirable default fxApp packages of the network operator.	fxManager Portal Service
33	Following a user to understand where it faces coverage issues or call drops. Once discovered, the system could send a message to the user acknowledging an issue and that the SP is trying to fix this. The information is uploaded to the cloud for further analysis.	Management {gatherStatistics}, General Framework, Messaging {sendSMS}
34	Reporting: This allows the developer to get reports and statistics of the fxDevice traffics, errors, and neighboring environment. The reporting criteria could be set up as: 1) Period, 2) Event based (reports only if a threshold is passed; a threshold could be based on absolute numbers or deltas), 3) One---time (i.e. the information).	Management {setMonitoring}, General Framework
35	Proprietary security protocol between the application in the mobile device and the BS in enterprise deployment. This would allow SP to create enterprise---class LTE for indoor using small cells.	Security {encryptContent, authenticateUser, Integrity, configureFirewall}, DataPath {RouteLocal}, General Framework
36	The femto cell at home that can't scan neighboring cells, gets a list of neighboring cells from the fxCloud or neighbors to broadcast to improve outbound handoffs.	Wireless {discoverNeighbors, retrieveNeighborInfo, setNeighborCells}, Management {retrieveNeighborInfo}
37	IMS---capable femto enabling various SP services to initiate or transferred between home network and mobile devices.	Messaging {sendfxCloud}, General Framework
38	Integrated smart home with multi---mode femto/WiFi router. For example, the application in the smart home router (fxDevice) allows for coordination of home physical security service (such as ADT services). With the fxCloudApp, the user could control the security status of his/her home and remotely control the setting.	Wireless {homeAreaNetworking}, Messaging {fxCloud}, General Framework
39	Home CCTV DVR implementation in an integrated home router. The camera feeds are recorded on the same router caching engine.	Cache {cacheVideo}, Messaging {sendfxCloud}, General Framework
40	Enterprise CCTV DVR implementation in an integrated router with cloud backup. The camera feeds are recorded in the cloud.	Cache {CacheVideo, CacheCloud}, Messaging {fxCloud}, General Framework
41	User traffic separation into various MPLS/LSP (integrating VPN and BS). In this case, users are classified into groups and each group is tunneled to the core network. For example, law enforcement use the encrypted tunnel while the other users use the defaulttunnel.	DataPath {Match, SetupTunnel, RouteTunnel}
42	Dynamic spectrum allocation where a BS borrows unused spectrum from the neighboring BS in case of temporary congestions.	DataPath {detectCongestion}, Wireless {discoverNeighbors, setWirelessConfig}, Management {retrieveNeighborInfo, borrowCarrier}, Messaging {fxCloud}

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0131

43	Proxy ANDSF (Access Network Discovery and Control Function) in the BS. When the UE tries to access the ANDSF, the Proxy ANDSF fxDeviceApp in the BS intercept the ANDSF messages and replies on behalf of the ANDSF. This would speed up the WiFi network discovery for example.	DataPath {matchTraffic}, General Framework, Messaging {sendfxCloud}
44	HD Audio calling: audio codecs in BS allows the phones with the right capabilities to communicate with HD Audio quality calling. If one of the mobile devices can't handle the HD quality audio, the BS performs the transcoding of the audio.	DataPath {matchTraffic}, General Framework
45	Enabling mobility on White Space Spectrum: The latest white space database will be downloaded. The mobile devices with special software applications make requests to the BS to retrieve the local available white space frequencies. This would allow the mobile device to handoff seamless between BS using the white space spectrum. The BS would advertise such a capability to the UEs. The secure communication allows peer---to---peer communication between the UE and the BS for certain applications such as White Space Channel lookup or site acquisition tools for field engineering and RF planning.	Messaging {sendfxCloud}, General Framework
46	Site acquisition tool where the technician uses a software tool on his mobile phone to interact with BSs to determine the best place to deploy the new BS. The phone app interests with fxDeviceApps in BSs and fxCloudApp. For example, fxDeviceApp performs air interface quality and gives a priority to traffic generated.	Messaging {sendfxCloud}, General Framework. DataPath {matchTraffic, setPriority}
47	DNS caching at BS to accelerate DNS lookup for subsequent users. This would improve the user experience by reducing response delays.	DataPath {matchTraffic}, General Framework
48	Continuous SLA monitory for certain protocols, applications or general link quality for jitter, delays, bandwidth.	DataPath {matchTraffic}, Management {gatherStatistics}, General Framework
49	Enterprise users can use the small cell deployed in their enterprise free of charge. The DPI engine in fxDevice classifies the traffic and identifies the enterprise private data and they would be routed locally and charging records are generated by the BS (the SP decides how to bill the enterprise user and it may decide to make such usage.	DataPath {matchTraffic, routeLocal}
50	On---premise small---cell aggregation where on fxDevice takes the responsibility of aggregating multiple small---cells on the premise and presenting them as a single BS/Node---B/eNodeB to the core network. In this case the aggregating fxDevice (AfxE) may take responsibility of local radio resource manager (RRM) or self---optimizing network (SON) server. AfxE can also implement and enforce policies concerning local service access on the premise (enterprise) as well as enable seamless handover between cellular technologies and on---premise (enterprise) local area network (LAN) for example using WiFi technology.	General Frameworks, Wireless {discoverNeighbors, retrieveNeighborInfo, setNeighborCells, setWirelessConfig}, Messaging {sendIntraApp}
51	Self adjusting fxCloud compute resources: The fxCloud creates and tears down virtual machines (VMs) based on the traffic load measured by itself or the packet processing and data plane engine. It is also possible that the VM adjustments are made based on the rush hour or pre---configured hours of the day or manually by the administrator.	CloudManagement {removeVM, createVM, detectCloudCongestion, actiavteVM, suspendVM}
52	Core Network Function Virtualization (cNFV) using distributed Software Defined Networking (dSDN): In the case, once an fxDevice is commissioned, its fxDeviceApp0 (i.e. firmware) uses predefined DNS names to discover the rest of the core network (e.g. MME, S---GW, etc.). The core network elements could be created.	General Frameworks, Messaging {sendInterApp, sendIntraApp), CloudManagement {deleteVM, suspendVM, createVM, activateVM, detectCloudCongestion}
53	Smart display and digital signage as wireless base stations: this use case allows combining the capabilities of a base station and the smart display in shopping malls, airports, and enterprises. There is a already electric power and connectivity for the smart display that the base station function could use.	General Frameworks
54	Tracking a mobile user at base station level even if the UE is in the idle mode. Currently in the idle mode, the MME, SGSN, or VLR/MSC can only know the location of the UE in with a accuracy of LA/RA/URA which is a very large area. The administrator uses the fxCloudApp and requests tracking of the user/UE (the admin may provide target area where the user/UE may be located; e.g. zip code or town name). The fxCloudApp contacts its fxDeviceApp in the fxDevices in the target area (TA) requesting information of a particular IMSI/MSISD/TMSI. Once an fxDevice responds indicating knowledge of the user location, the fxCloudApp directs that serving fxDevice to inform the fxDeviceApp in the next fxDevice of the instructions to track the user and contact fxCloudApp when they user hands off to due to its mobility. fxCloud could create user friendly interfaces and presentations to show for example the direction the user is moving, predict next location of the user, and estimate arrival of the user to a certain location.	DataPath {matchTraffic}, Messaging {sendIntraApp, sendfxCloud}
55	In case of emergency, fxDevices in a target area block all public traffic and only allow law enforcement sessions using specially encryption and authentication between the UE and the fxDevice or fxCloud (depending on the requirements).	DataPath {matchTraffic, setPriority}, Security {configureFirewall, encryptTraffic, authenticateUser}, General Framework

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0132

56	SON triggered by the data traffic consumption. In this case, the mobile cell configuration is adjusted based on the current traffic of the cell. It could for example force handoff some users (high users or users moving fast) to neighboring cells.	DataPath {identifyTopDevices}, Wireless {forceHandoff}
57	Time based QoS and policy. In this use case, the infManager pushes certain policies in target area (TA). Such polices could be turning on extra capacity (frequency channels) or QoS policies and traffic prioritization on the WAN like.	DataPath {matchTraffic, setPriority}, Wireless {discoverNeighbors, setWirelessConfig}
58	Fast moving users detection: in this use case, the application in the BS measures the speed of the user (based on the time was handed off to this cell until it was handed over to the next cell). This information could be used by the fxDevice or fxCloud to force handoff the mobile user to a larger cell to reduce handoff rates.	General Frameworks, Wireless {identifyTopMovers, forceHandoff}
59	Emergency broadcast of information and live camera feeds. In case of emergency, the enterprise network administrator temporarily block all traffic and only allow emergency related traffic across the target area in the enterprise network. For example, the network admin broadcasts the camera feeds and alerts over the WiFi access points using multicast IP techniques. The client devices and digital sinages can listen to that particular multicast and show the camera feeds and emergency alerts.	DataPath {matchTraffic, setPriority}, Security {configureFirewall}, General Framework
60	Dynamic advertising in shopping malls. The shopper would use an app from the shopping mall. When the shopper brings up or refreshes the app, the dSDN allows the shopping mall to pinpoint the exact indoor/outdoor location of the user (e.g. the shop the shopper is in) and send special coupon and ads for that store to the shopper.	General Frameworks, Wireless {positionDevice}, Messaging {sendIntraApp, sendfxCloud}, Security {configureFirewall}, DataPath {matchTraffic, redirectTraffic}
61	Smart tour guide. The tourist would use an app from the tourism board or museum. When the user brings up or refreshes the app, the dSDN allows to pinpoint the exact location of the tourist (e.g. the painting he/she is looking at) and send tour guide information (e.g. audio, video, hypertext). dSDN allows to localize the positioning.	General Frameworks, Wireless {positionDevice}, Messaging {sendIntraApp, sendfxCloud}, Security {configureFirewall}, DataPath {matchTraffic, redirectTraffic}
62	General hotspot landing page	DataPath {matchTraffic, redirectTraffic}, Security {authenticateUser}
63	Proprietary wireless mesh implementation	Wireless {discoverNeighbor, retrieveNeighborInfo, measureInterference, setWirelessConfig}, General Framework, Messaging {sendIntraApp}
64	Power and energy consumption map of the network. A real-–time fxCloudApp pulls the power consumption details from each fxDevice periodicaly or on demand. The pulled data then are combined with the location information and correlated into a map to present a viewable consumption map.	Management {getfxDeviceLocation, getPowerConsumptionRate}
65	Network deployment optimization. This application allows network operator to get advice for best optimal reconfiguration and/or redeployment of the network. For example, the network admin could get suggestion that based on the power consumptions and coverage perceived and spectrum utilization, it might be better to replace a particular group of smaller cells with a single macro-–cell. This app could also pull information from larger OEMs based on the subject deployment scenario (the OEMs either provide this information statically to the tool or provide network APIs for use by this app e.g. using REST protocol).	Management {getfxDeviceLocation, getPowerConsumptionRate, gatherStatistics}
66	High availability. This app allows two or more fxDevice to form a redundancy cluster where they participate in statueful redundancy and load-–balancing.	Messaging {sendIntraApp}, General Framework
67	Bandwidth Bursting. This feature enables the dynamic adjustment of backhaul or WAN link bandwidth as the demands change in the enterprise or public wireless deployments. In this case, the fxDeviceApp in the fxDevice monitors the WAN link utilization, once a preset threshold percentage has reached, it messages its sister fxCloudApp which in turn would send a request to the fxDeviceApp to increase or decrease the WAN bandwidth. If additional backend reconfiguration is needed, the fxCloudApp performs those additional changes prior to informing the fxDeviceApp of the bandwidth.	Management {getStatistics, configureWAN}
68	Wireless link quality measurement by the small-–cell or WiFi AP: The wireless SoC in the small-–cell or WiFi AP continuously monitors the link quality for power adjustment and handover reasons. Such information could be exposed to the app developers to create cell-–level analytics and reports without the need to have test-– drives or field testing of service quality.	Management {gatherStatistics}, General Framework
69	Data Path Consolidation: An application on the fxDevice discovers identical data feeds from a cloud server to two or more users. The fxDeviceApp in the fxDevice can discover this potential duplicate flows at the session establishment (via a proxy function) and it would consolidate those sessions into one on the uplink towards the cloud server. A good example of this use case is the web-–based desktop sharing where different many users in the same office site get the same feed from the cloud server.	FastPath {detectIdenticalFlows, aggregateFlows}

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0133

[0117] The following references are herein specifically incorporated by reference in their entirety into this disclosure:

[1] 3GPP TS 23.002, "Network Architecture";

[2] 3GPP TS 23.402, "Architecture enhancements for non-3GPP accesses";

[3] 3GPP TS 23.401, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access";

[4] Open Base Station Architecture Initiative (OBSAI); and

[5] SPDY Protocol (http://datatracker.ietf.org/doc/draft-mbelshe-httpbis-spdy/).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0134

CLAIMS

1. A system for processing data packets in a network, the system comprising:

a programmable network device hosting a plurality of first network applications;

a programmable cloud device hosting a plurality of second network applications;

wherein at least one of the plurality of first network applications in the programmable network device and at least one of the plurality of second network applications in the programmable cloud device are in secure communication with each other to form a distributed application;

the programmable network device and programmable cloud device are each powered by a sandboxing operating system which facilitates deployment of the plurality of first and second network applications and facilitates upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device;

wherein the programmable network device verifies authenticity of the upgrades to the plurality of first network applications and the programmable cloud device verifies authenticity of the upgrades to the plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

an application repository storing distributed applications which have been tested for installation in the programmable network device and programmable cloud device; and

an application management portal coupled to the programmable network device, programmable cloud device and the application repository, wherein the application management portal manages usage of the distributed applications on the programmable network device and programmable cloud device.

2. The system of claim 1, wherein each of the programmable network device and programmable cloud device include an access network interface unit to send and receive communications and a processor with a memory associated with the network interface unit.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0135

3. The system of claim 1, wherein the at least one of the first network applications is responsible for a plurality of functional capabilities of the programmable network device and is capable of at least one of the group consisting of: updating the firmware of the programmable network device; updating default communication protocols of the programmable network device; initializing programmable network device settings; boot loading the programmable network device; providing security digital certificates; providing operating system (OS) and kernel upgrades; and upgrading and commissioning agents for upgrades of the at least one of the first network applications.

4. The system of claim 1, wherein at least one of the second network applications is responsible for a plurality of functional capabilities of the programmable cloud device and is capable of at least one of the group consisting of: upgrading and commissioning agents for upgrades of the at least one of the second network applications and related procedures; hypervisor upgrades; virtual machine and isolated execution environment lifecycle management; security digital certificate management; operating system (OS) and kernel upgrades; and default communication protocols.

5. The system of claim 1, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines, the network functions consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).

6. The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a virtual fabric which provides a secure communication layer for at the least one of the plurality of first network applications and the at least one of the plurality of second network applications.

7. The system of claim 1, further comprising:

a second programmable network device which includes the at least one of the

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0136

plurality of first network applications which is in the first programmable network device.

8. The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a distributed resource service (DRS) which controls access to a plurality of resources by the plurality of first network applications and the plurality of the second network applications upon receiving instructions from the application management portal.

9. The system of claim 1, wherein the plurality of resources include at least one of a group consisting of: a firewall; data and statistics; storage; compute; load evaluator; general settings; routing engine; virtual machine master controller; power controller; and wireless engine.

10. The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a distributed resource service (DRS) which is capable of providing at least one service from a group consisting of: exposing application programming interfaces (APIs) to other applications; configuring and managing platform resources; policy enforcement and authorization of applications access to platform resource and other application's APIs; and policy conflict resolution.

11. The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a distributed content service wherein contents generated by the at least one of the plurality of first network applications may be shared with the at least one of a plurality of second network applications.

12. The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a distributed notification service wherein notice of a particular event is sent from at least one of the plurality of first network applications to at least one of the plurality of

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0137

second network applications.

13. The system of claim 10, wherein the DRS further includes a load controller; and wherein the load controller is adapted to monitor loads on at least one of the plurality of first network applications and at least one of the plurality of second network applications and effect change in accordance with thresholds received from the application management portal.

14. The system of claim 1, wherein the programmable network device, at least one of the plurality of first network applications, the programmable cloud device, and at least one of the plurality of second network applications each have unique identities.

15. A system for processing data packets, the system comprising:

a programmable network device processing data flows hosting at least one of a plurality of first network applications;

a programmable cloud device having a plurality of virtual machines wherein at least one of the virtual machines hosts at least one of a plurality of second network applications;

wherein the at least one of the plurality of first and the at least one of the plurality of second network applications are in secure communication to form a distributed application;

the programmable network device and programmable cloud device are each powered by a sandboxing operation system which facilitates deployment of the plurality of first and second network applications and facilitates upgrade of the plurality of first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device;

wherein the programmable network device verifies authenticity of the upgrades to the plurality of first network applications and the programmable cloud device verifies the authenticity of the upgrades to the plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

an application repository storing distributed applications which have been tested

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0138

for installation in the programmable network device and programmable cloud device; and

an application management portal coupled to the programmable network device, programmable cloud device and the application repository, wherein the application management portal manages provisioning, usage and de-provisioning of the distributed applications on the programmable network device and programmable cloud device.

16. The system of claim 15, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines, the network functions consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).

17. A system for processing data packets in a network, the system comprising:-

a programmable network device processing data flows, hosting at least one of a plurality of first network applications;

a programmable cloud device partitioned into a plurality of zones, wherein at least one of the plurality of zones hosts at least one of a plurality of ~~first~~ second network applications;

wherein the at least one of the plurality of first network applications and at least one of the plurality of second network applications are in secure communication to form a distributed application;

the programmable network device and programmable cloud device are each powered by a sandboxing operation system which facilitates deployment of the plurality of first and second network applications and facilitates upgrade of the plurality of first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device;

wherein the programmable network device verifies authenticity and integrity of upgrades to the plurality of first network applications and the programmable cloud device verifies the authenticity of the upgrades to the plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0139

an application repository storing distributed applications which have been tested for installation in the programmable network device and programmable cloud device; and

an application management portal coupled to the programmable network device, programmable cloud device and the application repository, wherein the application management portal manages usage of the distributed applications on the programmable network device and programmable cloud device.

18. The system of claim 17, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines, the network functions consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0140

ABSTRACT

The Distributed Software Defined Network (dSDN) disclosed herein is an end-to-end architecture that enables secure and flexible programmability across a network with full lifecycle management of services and infrastructure applications (fxDeviceApp). The dSDN also harmonizes application deployment across the network independent of the hardware vendor. As a result, the dSDN simplifies the network deployment lifecycle from concept to design to implementation to decommissioning.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0141

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0142

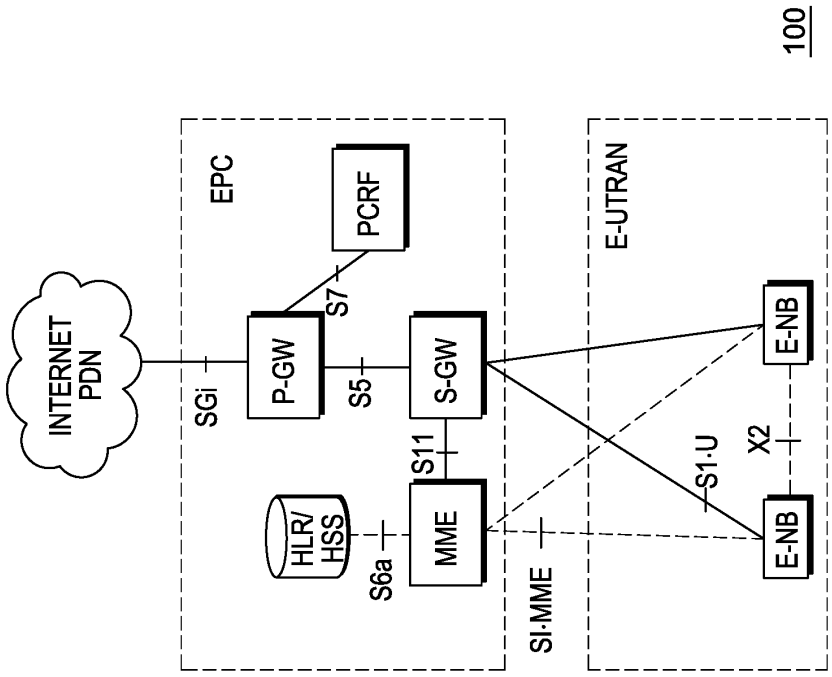


FIG. 1

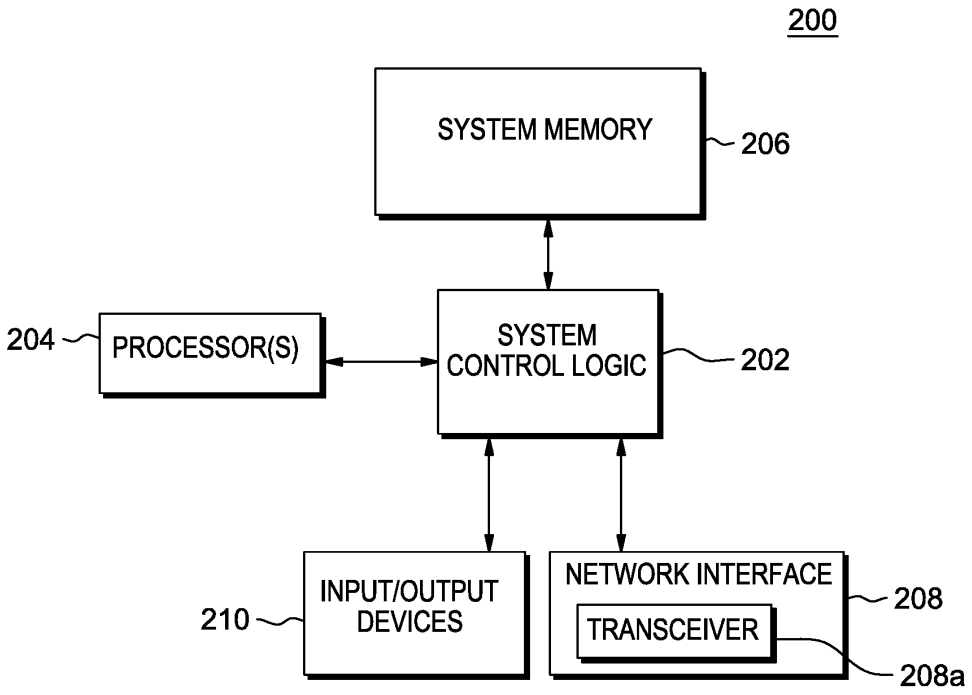

200
SYSTEM MEMORY
206
204
PROCESSOR(S)
SYSTEM
CONTROL LOGIC
202
210
INPUT/OUTPUT
DEVICES
NETWORK INTERFACE
208
TRANSCEIVER
208a

FIG. 2

FIG. 3

Docket No. TAAG001CONT1

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0144

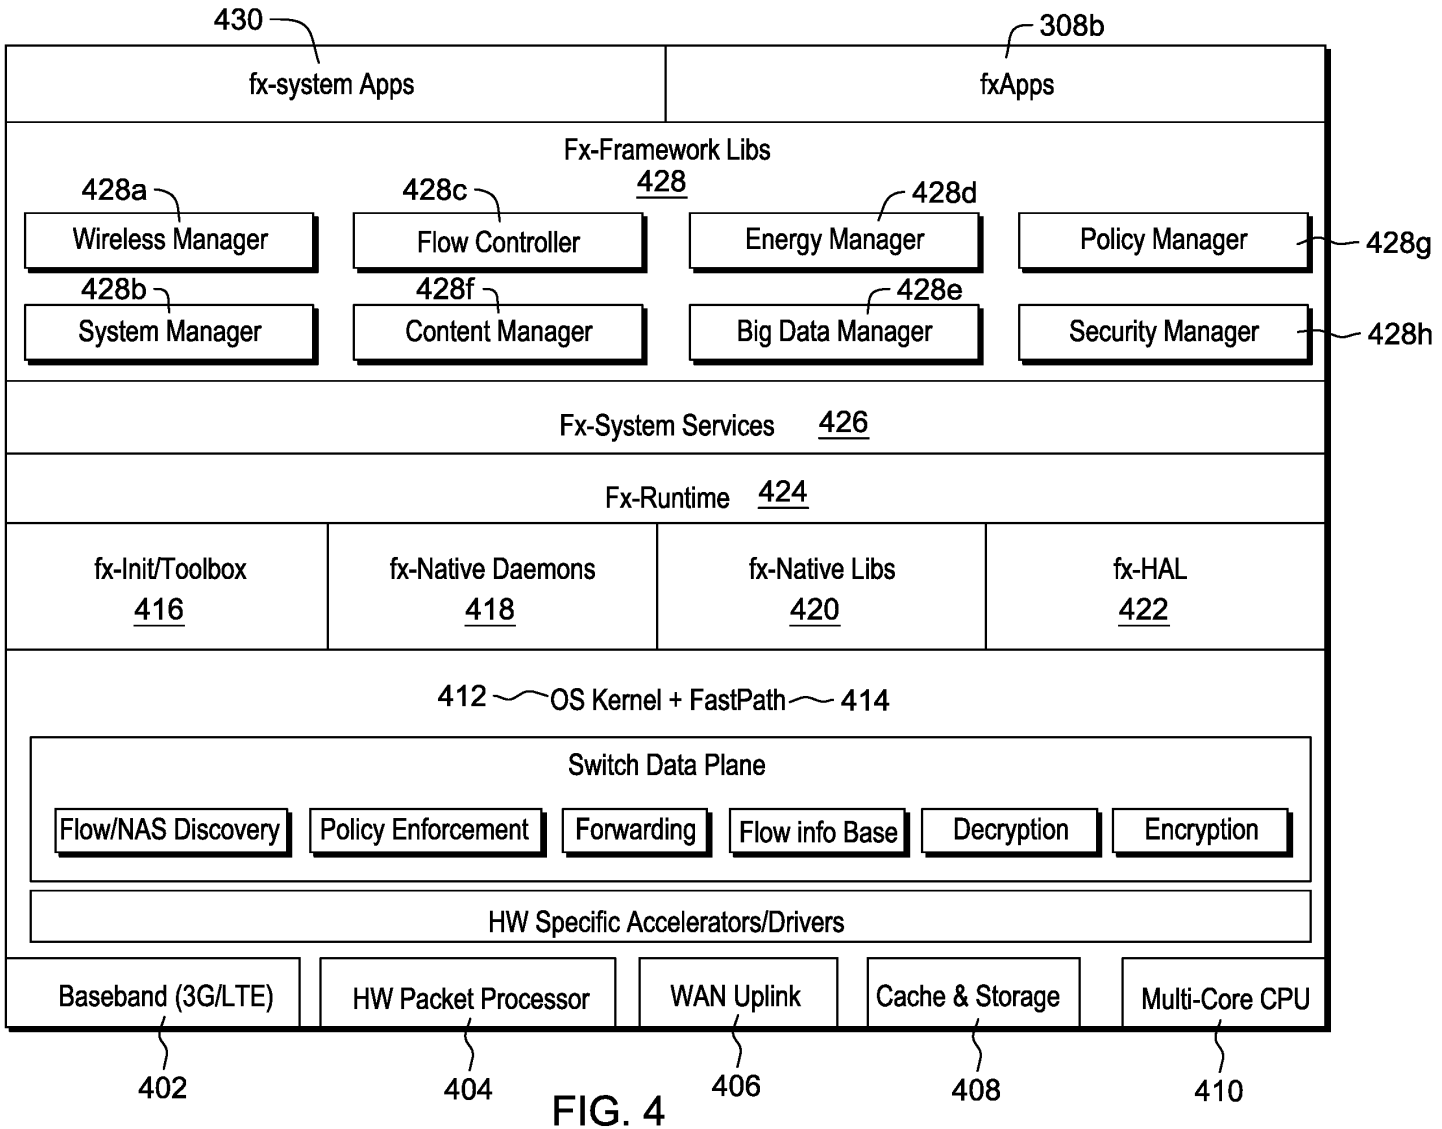

430
308b
fx-system Apps
fxApps
Fx-Framework Libs
428
428a
Wireless Manager
428c
Flow Controller
428d
Energy Manager
Policy Manager
428g
428b
System Manager
428f
Content Manager
428e
Big Data Manager
Security Manager
428h
Fx-System Services 426
Fx-Runtime 424
fx-Init/Toolbox
416
fx-Native Daemons
418
fx-Native Libs
420
fx-HAL
422
412
OS Kernel + FastPath
414
Switch Data Plane
Flow/NAS Discovery
Policy Enforcement
Forwarding
Flow info Base
Decryption
Encryption
HW Specific Accelerators/Drivers
Baseband (3G/LTE)
HW Packet Processor
WAN Uplink
Cache & Storage
Multi-Core CPU
402
404
406
408
410

FIG. 4

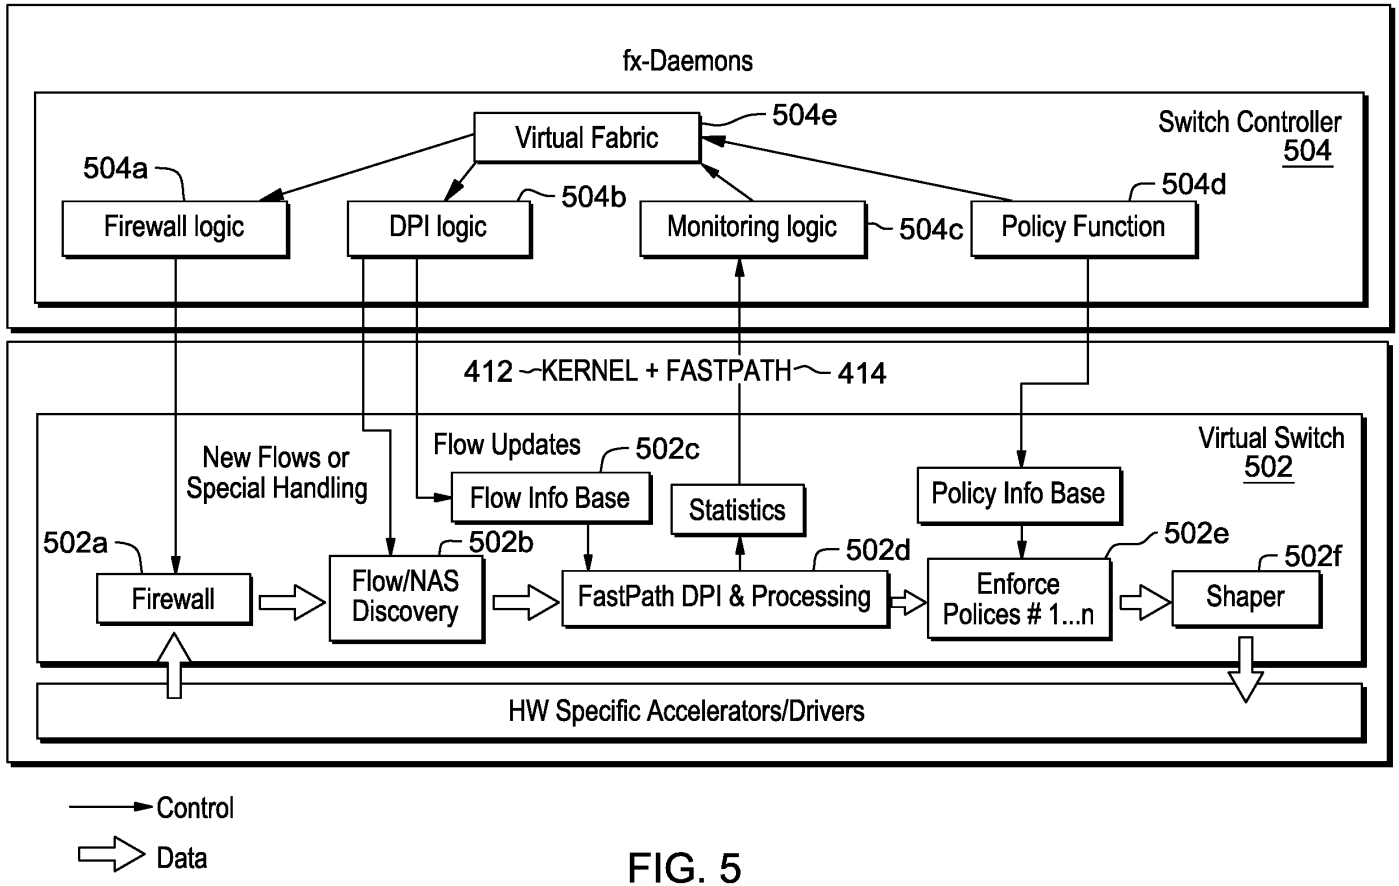

fx-Daemons
Virtual Fabric
504e
Switch Controller
504
504a
Firewall logic
DPI logic
504b
Monitoring logic
504c
Policy Function
504d
412
KERNEL + FASTPATH
414
New Flows or
Special Handling
Flow Updates
502c
Flow Info Base
Statistics
Policy Info Base
Virtual Switch
502
502a
Firewall
Flow/NAS
Discovery
502b
FastPath DPI & Processing
502d
Enforce
Polices # 1...n
502e
Shaper
502f
HW Specific Accelerators/Drivers
Control
Data

FIG. 5

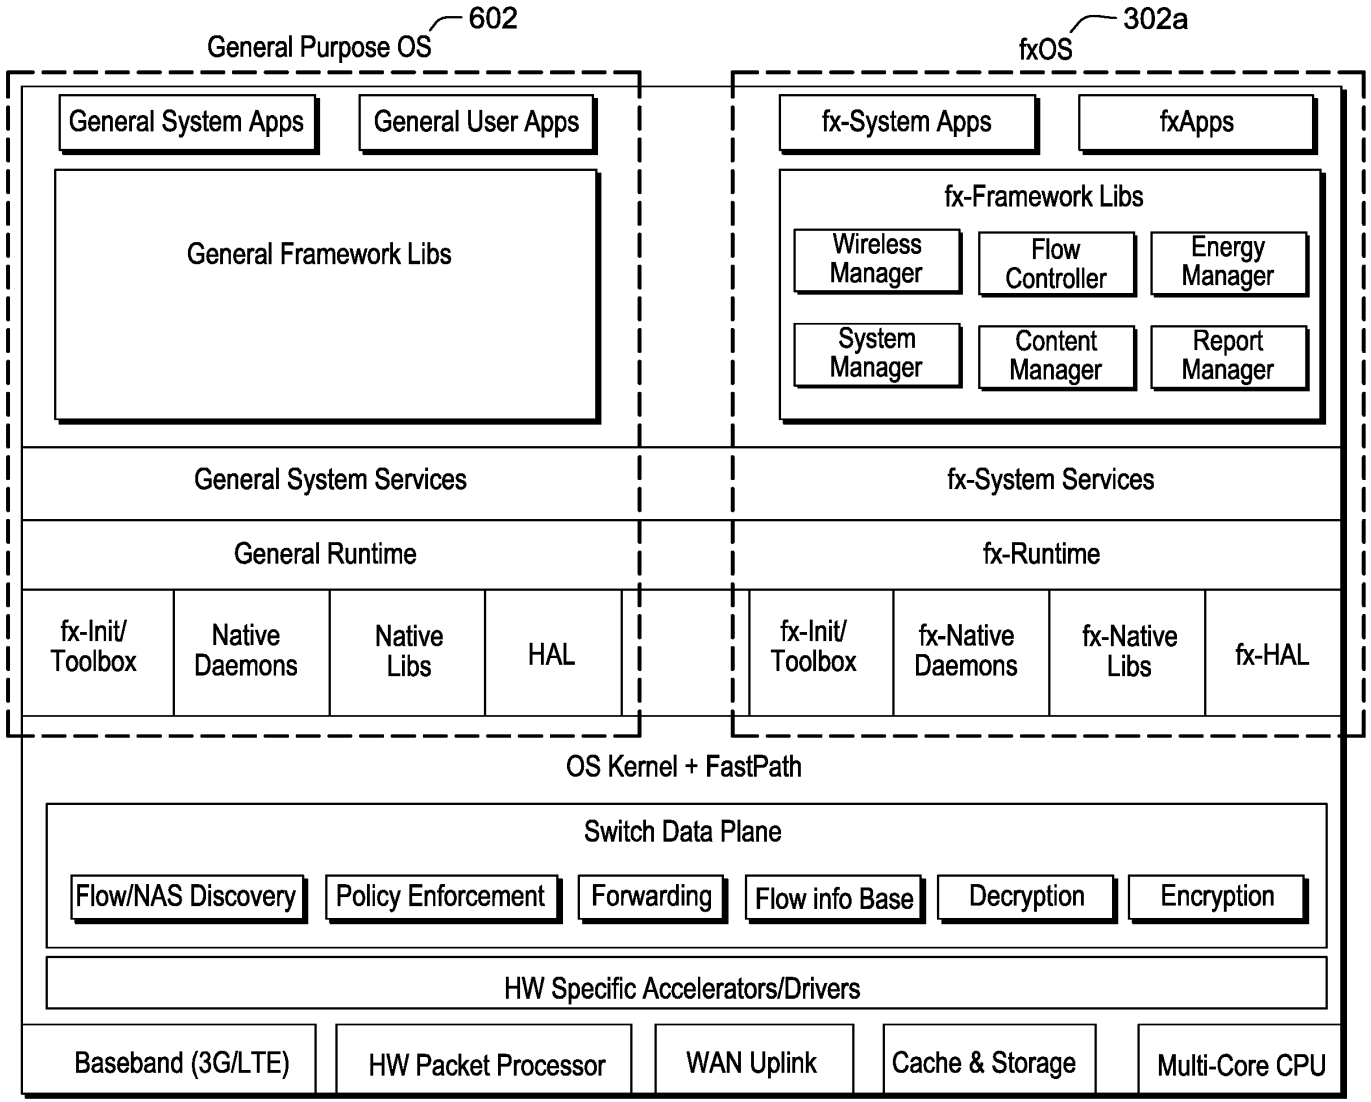

602
General Purpose OS
302a
fxOS
General System Apps
General User Apps
General Framework Libs
fx-System Apps
fxApps
fx-Framework Libs
Wireless Manager
Flow Controller
Energy Manager
System Manager
Content Manager
Report Manager
General System Services
fx-System Services
General Runtime
fx-Runtime
fx-Init/ Toolbox
Native Daemons
Native Libs
HAL
fx-Init/ Toolbox
fx-Native Daemons
fx-Native Libs
fx-HAL
OS Kernel + FastPath
Switch Data Plane
Flow/NAS Discovery
Policy Enforcement
Forwarding
Flow info Base
Decryption
Encryption
HW Specific Accelerators/Drivers
Baseband (3G/LTE)
HW Packet Processor
WAN Uplink
Cache & Storage
Multi-Core CPU

FIG. 6

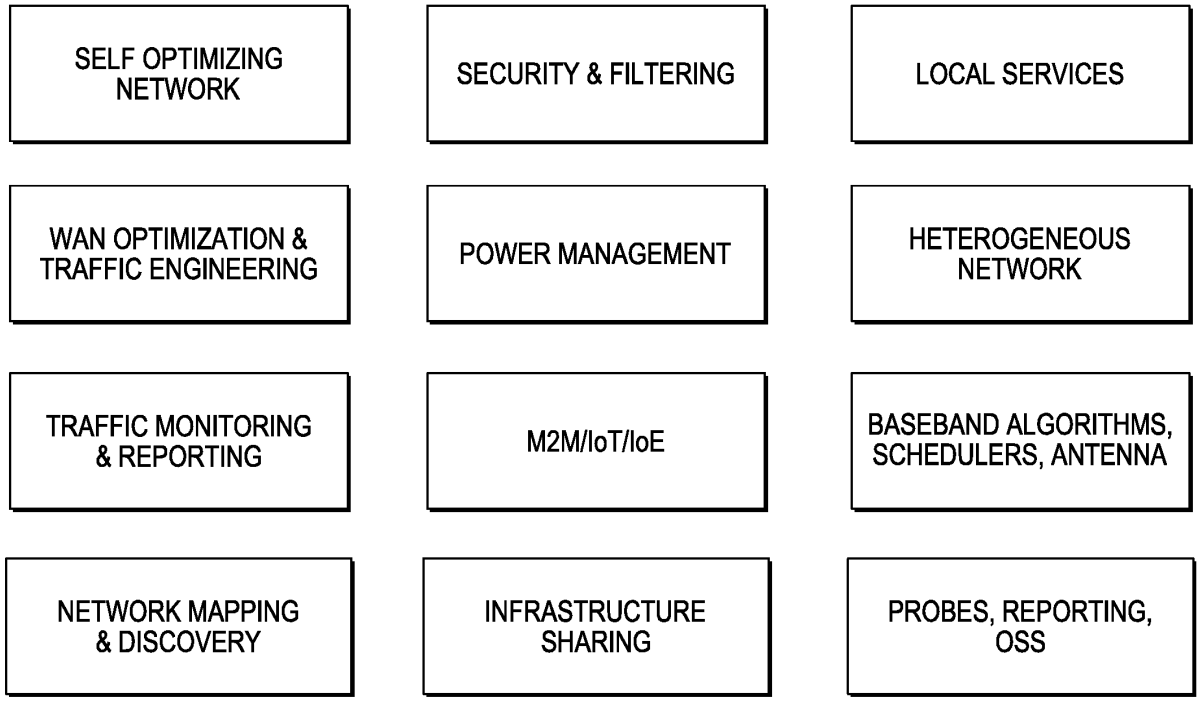

SELF OPTIMIZING NETWORK
SECURITY & FILTERING
LOCAL SERVICES
WAN OPTIMIZATION & TRAFFIC ENGINEERING
POWER MANAGEMENT
HETEROGENEOUS NETWORK
TRAFFIC MONITORING & REPORTING
M2M/IoT/IoE
BASEBAND ALGORITHMS, SCHEDULERS, ANTENNA
NETWORK MAPPING & DISCOVERY
INFRASTRUCTURE SHARING
PROBES, REPORTING, OSS

FIG. 7

FIG. 8

Docket No. TAAG001 CONT1

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0149

FIG. 9

Docket No. TAAG001 CONT1

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0150

FIG. 10

Docket No. TAAG001 CONT1

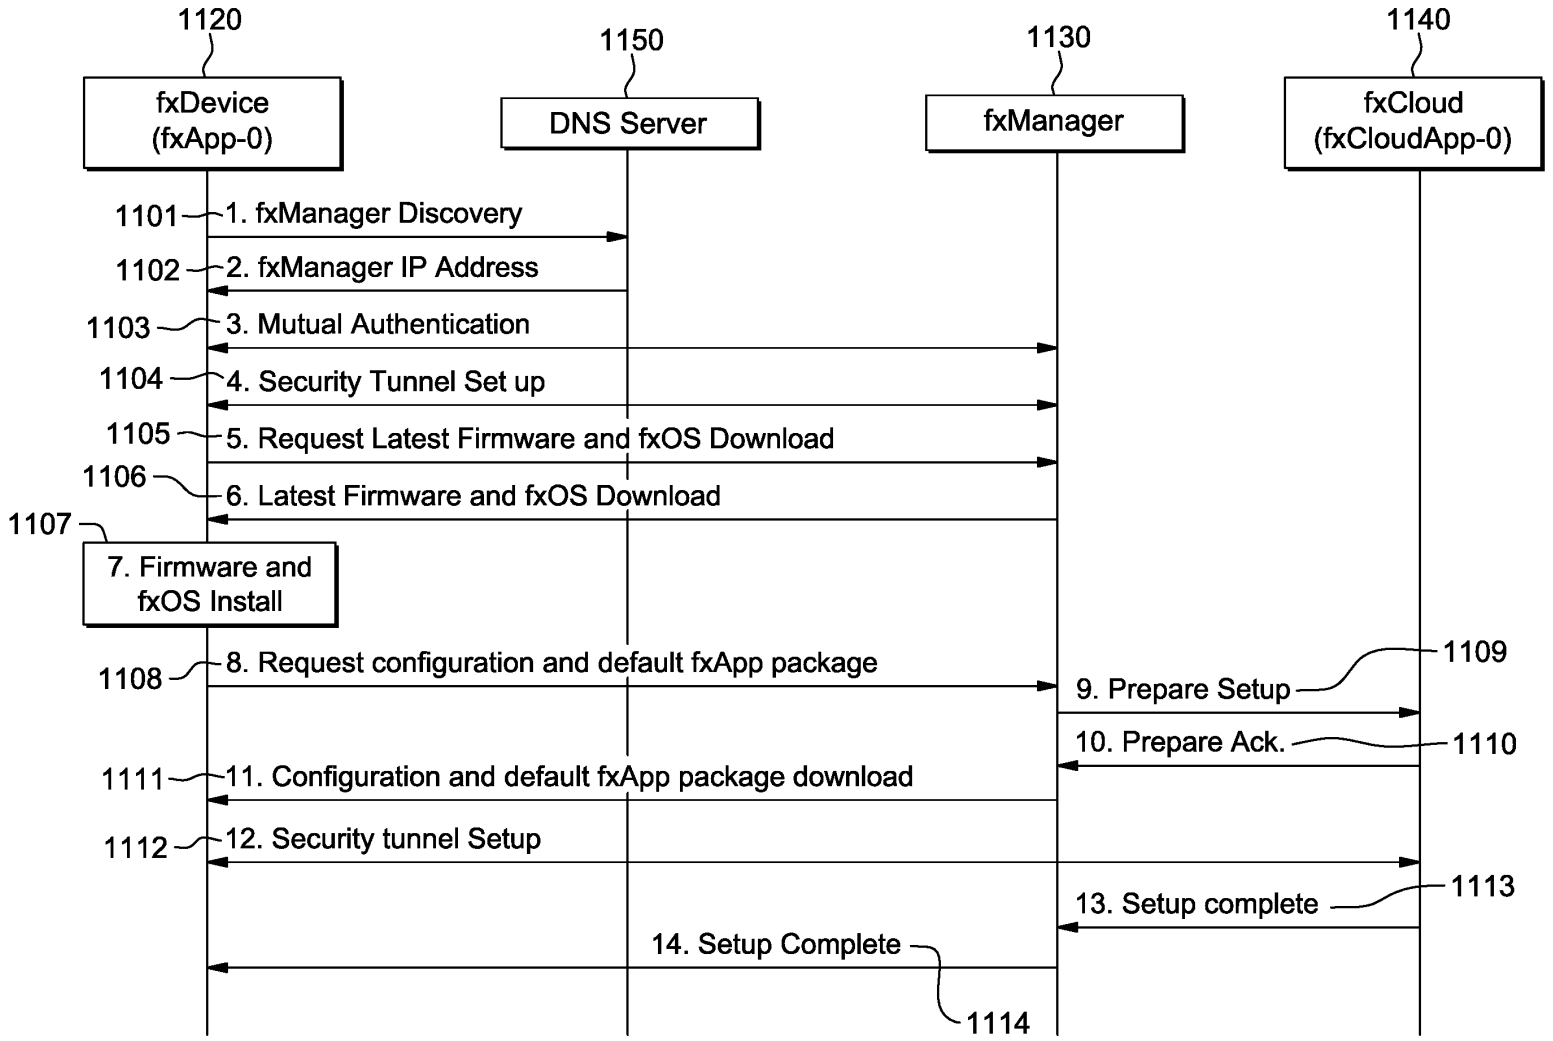

1120
fxDevice
(fxApp-0)
1150
DNS Server
1130
fxManager
1140
fxCloud
(fxCloudApp-0)
1101
1. fxManager Discovery
1102
2. fxManager IP Address
1103
3. Mutual Authentication
1104
4. Security Tunnel Set up
1105
5. Request Latest Firmware and fxOS Download
1106
6. Latest Firmware and fxOS Download
1107
7. Firmware and
fxOS Install
1108
8. Request configuration and default fxApp package
1109
9. Prepare Setup
10. Prepare Ack.
1110
1111
11. Configuration and default fxApp package download
1112
12. Security tunnel Setup
1113
13. Setup complete
14. Setup Complete
1114

FIG. 11

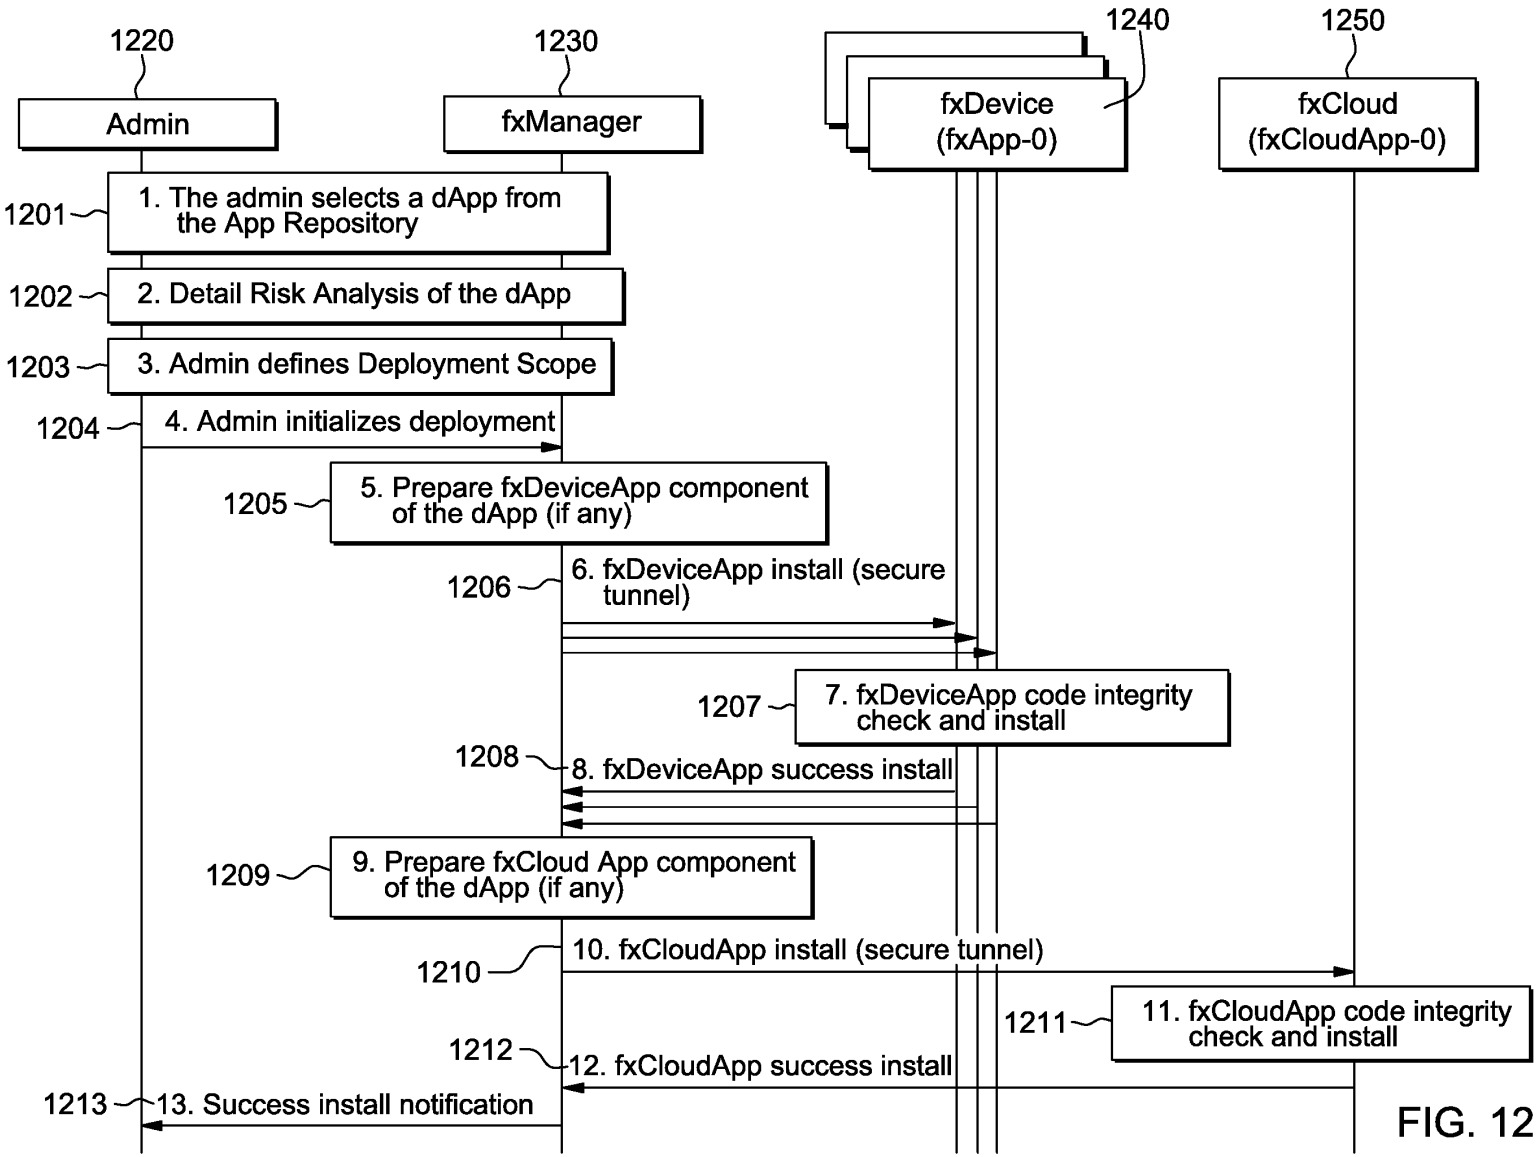

1220
Admin
1230
fxManager
1240
fxDevice
(fxApp-0)
1250
fxCloud
(fxCloudApp-0)
1201
1. The admin selects a dApp from the App Repository
1202
2. Detail Risk Analysis of the dApp
1203
3. Admin defines Deployment Scope
1204
4. Admin initializes deployment
1205
5. Prepare fxDeviceApp component of the dApp (if any)
1206
6. fxDeviceApp install (secure tunnel)
1207
7. fxDeviceApp code integrity check and install
1208
8. fxDeviceApp success install
1209
9. Prepare fxCloud App component of the dApp (if any)
1210
10. fxCloudApp install (secure tunnel)
1211
11. fxCloudApp code integrity check and install
1212
12. fxCloudApp success install
1213
13. Success install notification

FIG. 12

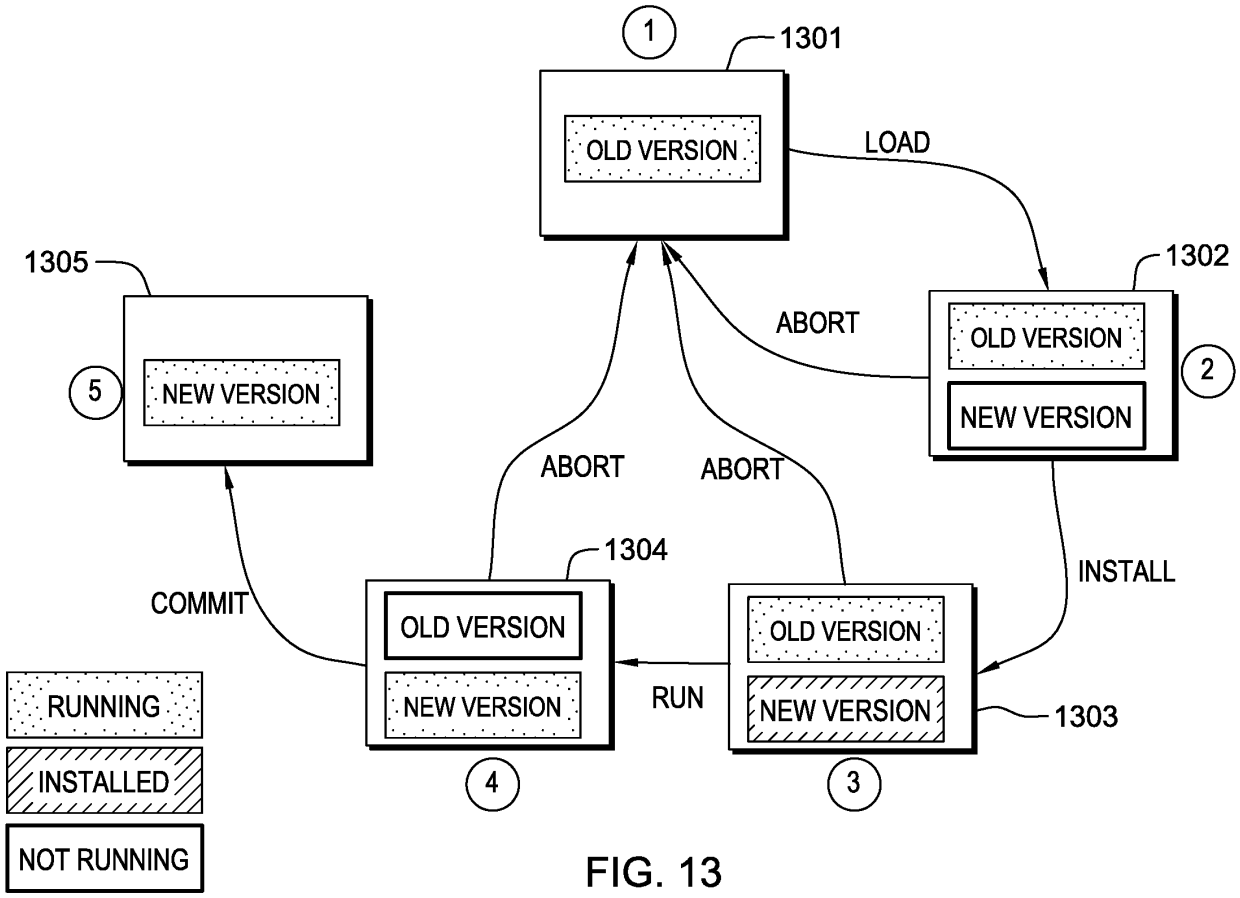

1
1301
OLD VERSION
LOAD
1305
1302
ABORT
OLD VERSION
5
NEW VERSION
NEW VERSION
2
ABORT
ABORT
1304
INSTALL
COMMIT
OLD VERSION
OLD VERSION
NEW VERSION
RUN
NEW VERSION
1303
RUNNING
INSTALLED
NOT RUNNING
4
3

FIG. 13

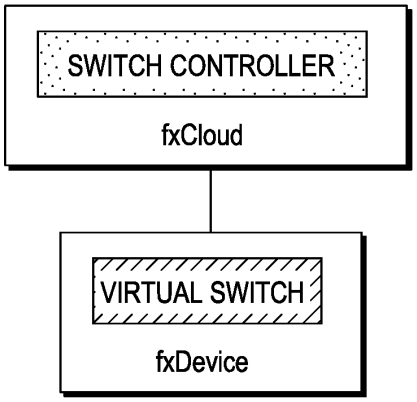


FIG. 14A

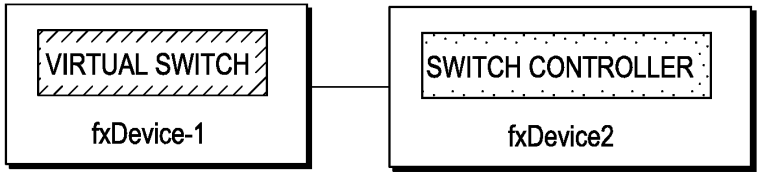


FIG. 14B

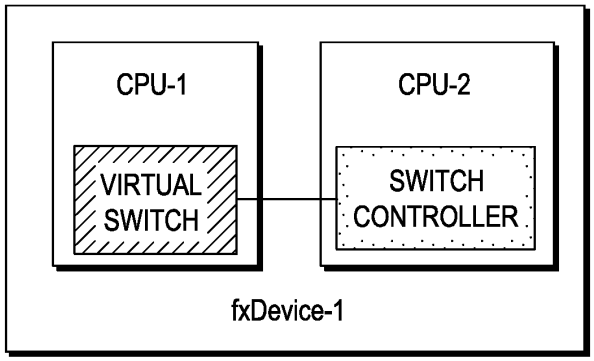


FIG. 14C

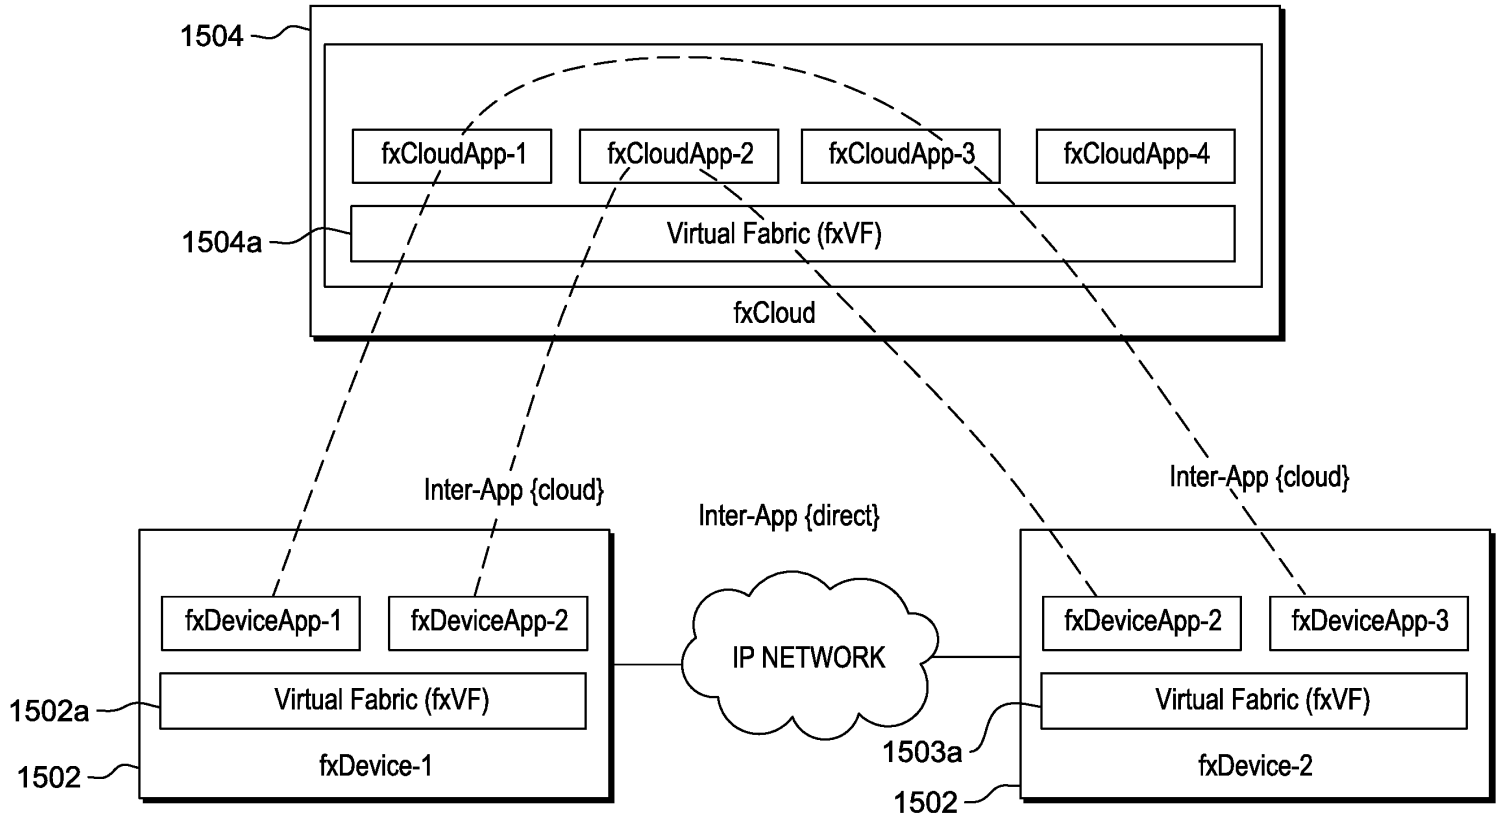

1504
fxCloudApp-1
fxCloudApp-2
fxCloudApp-3
fxCloudApp-4
1504a
Virtual Fabric (fxVF)
fxCloud
Inter-App {cloud}
Inter-App {direct}
Inter-App {cloud}
fxDeviceApp-1
fxDeviceApp-2
IP NETWORK
fxDeviceApp-2
fxDeviceApp-3
1502a
Virtual Fabric (fxVF)
Virtual Fabric (fxVF)
1503a
1502
fxDevice-1
fxDevice-2
1502

FIG. 15

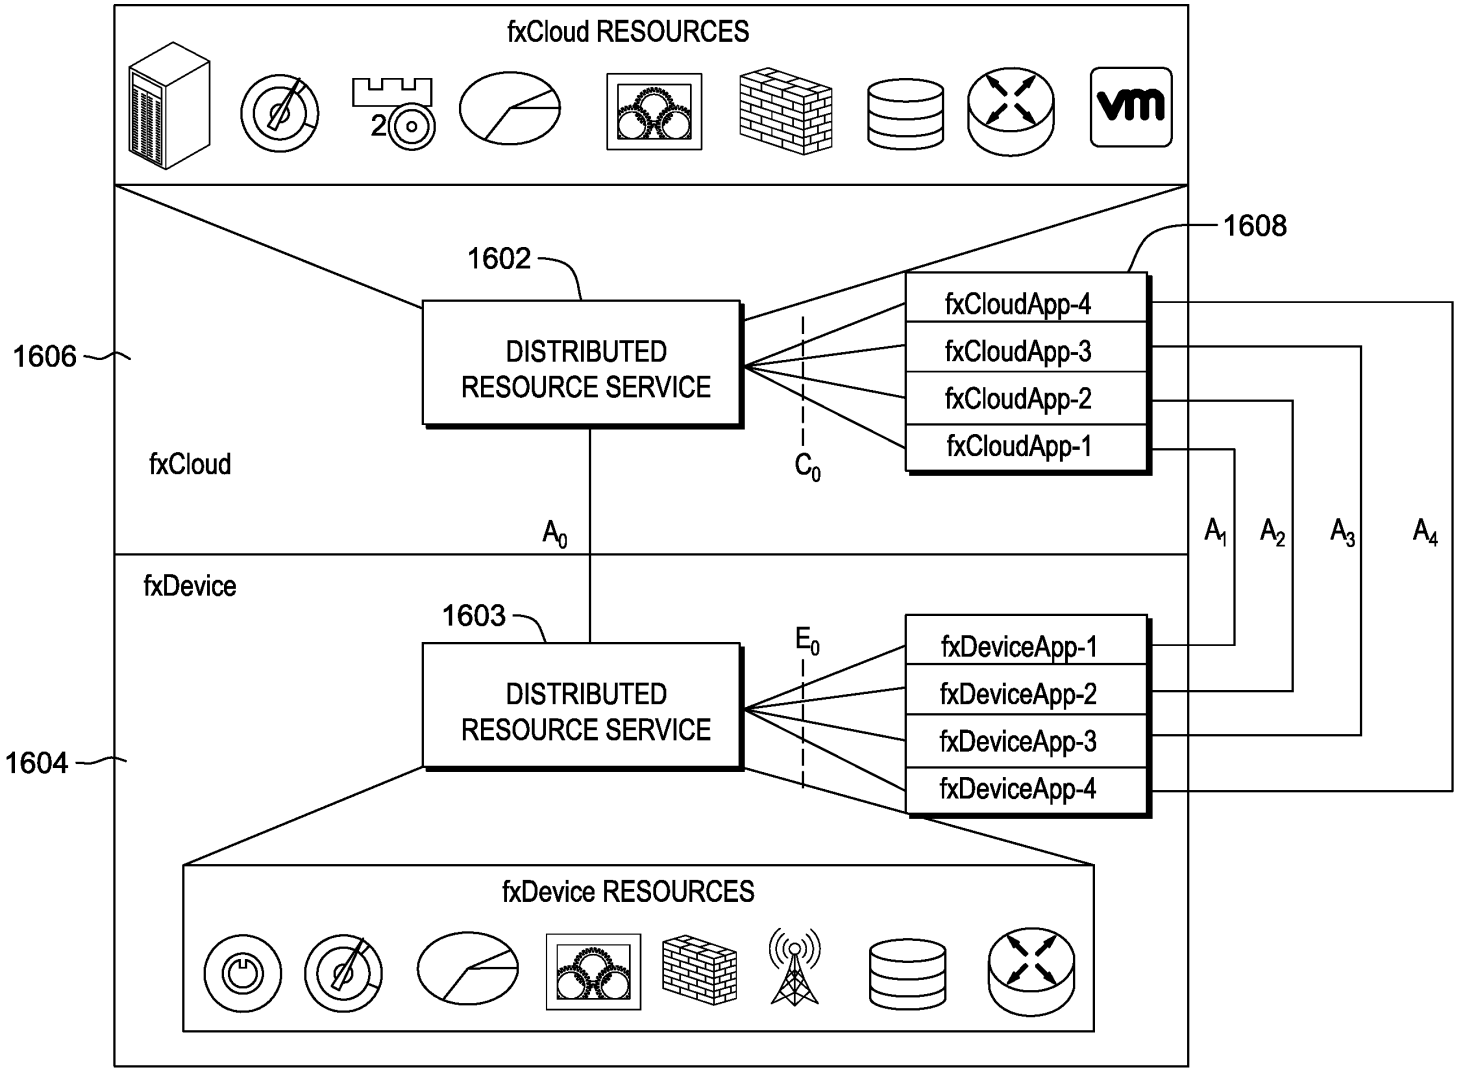

fxCloud RESOURCES
1602
1608
1606
DISTRIBUTED
RESOURCE SERVICE
fxCloudApp-4
fxCloudApp-3
fxCloudApp-2
fxCloudApp-1
fxCloud
C_0
A_0
A_1
A_2
A_3
A_4
fxDevice
1603
E_0
DISTRIBUTED
RESOURCE SERVICE
fxDeviceApp-1
fxDeviceApp-2
fxDeviceApp-3
fxDeviceApp-4
1604
fxDevice RESOURCES

FIG. 16

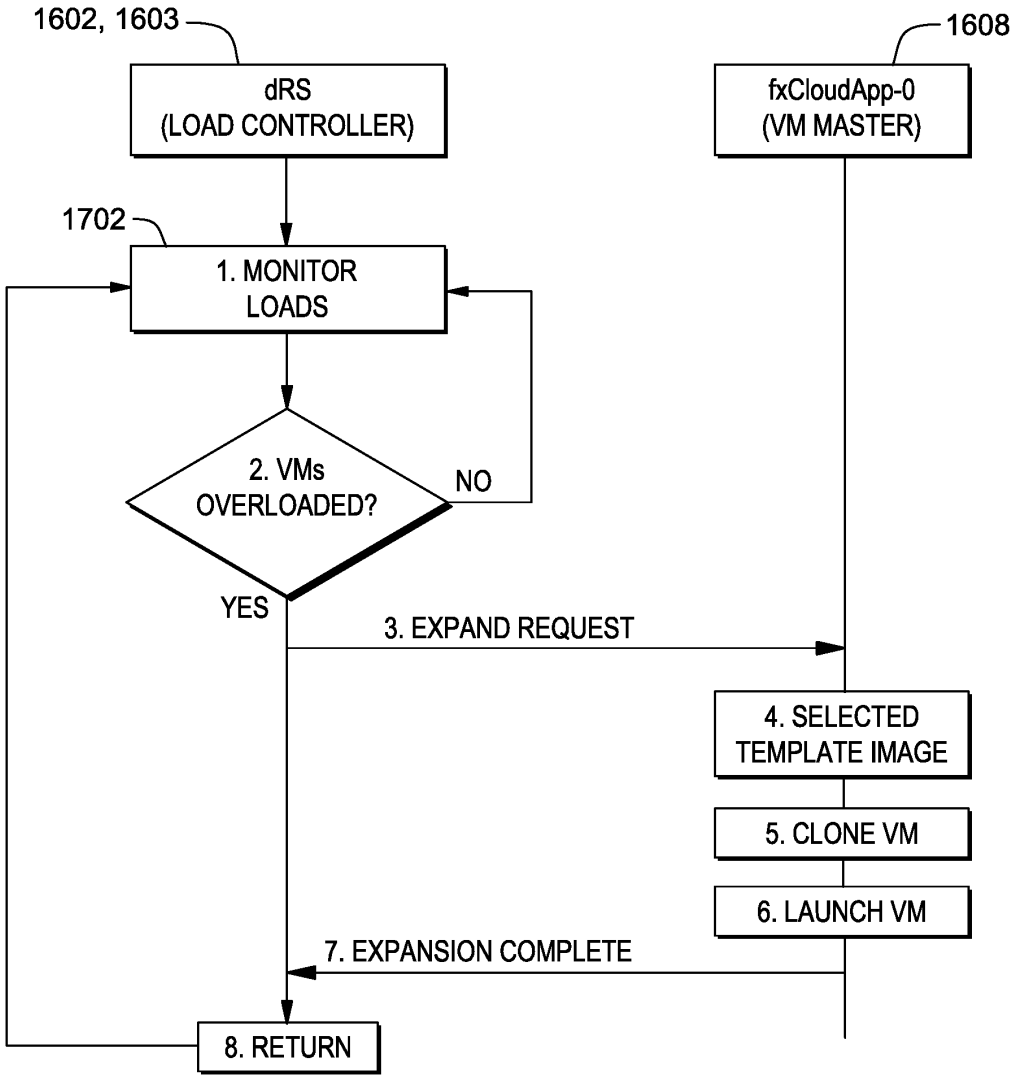

1602, 1603
dRS
(LOAD CONTROLLER)
1608
fxCloudApp-0
(VM MASTER)
1702
1. MONITOR
LOADS
2. VMs
OVERLOADED?
NO
YES
3. EXPAND REQUEST
4. SELECTED
TEMPLATE IMAGE
5. CLONE VM
6. LAUNCH VM
7. EXPANSION COMPLETE
8. RETURN

FIG. 17A

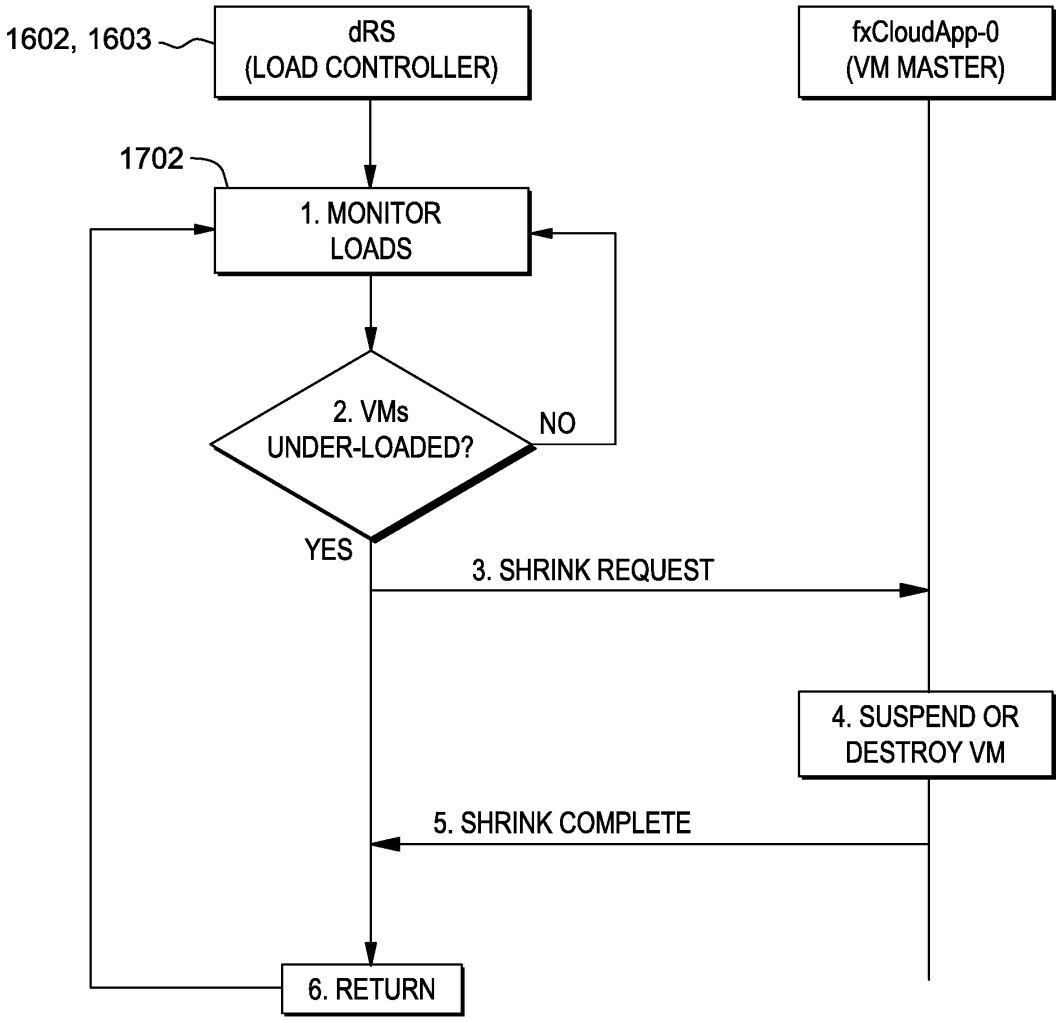

1602, 1603
dRS
(LOAD CONTROLLER)
fxCloudApp-0
(VM MASTER)
1702
1. MONITOR
LOADS
2. VMs
UNDER-LOADED?
NO
YES
3. SHRINK REQUEST
4. SUSPEND OR
DESTROY VM
5. SHRINK COMPLETE
6. RETURN

FIG. 17B

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Pouya		Taaghol	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	1260 Trestlewood Lane		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95138	Country	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Vivek		Ramanna	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	549 Manhattan Place		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95136	Country	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

☐ An Address is being provided for the correspondence Information of this application.			
Customer Number	112449		
Email Address	chris.horgan@roarkip.com	Add Email	Remove Email

Application Information:

Title of the Invention	Distributed Software Defined Networking		
Attorney Docket Number	TAAG001CONT1	**Small Entity Status Claimed**	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	5	**Suggested Figure for Publication (if any)**	1

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☒ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	112449		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	Continuation of	14295331	2014-06-04
Prior Application Status	Expired		Remove

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
14295331	non provisional of	61634807	2013-06-13

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)[i] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Remove

Application Number	Country[i]	Filing Date (YYYY-MM-DD)	Access Code[j] (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date o f filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

◯ Assignee	◯ Legal Representative under 35 U.S.C. 117	◯ Joint Inventor
◯ Person to whom the inventor is obligated to assign.	◯ Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. ☐

Prefix	Given Name	Middle Name	Family Name	Suffix

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Mailing Address Information For Applicant:			
Address 1			
Address 2			
City		**State/Province**	
Country		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Applicant Data may be generated within this form by selecting the Add button.

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

If the Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Non-Applicant Assignee:			
Address 1			
Address 2			
City		**State/Province**	
Country		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button.

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Signature:

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications.					
Signature	/Christopher J. Horgan/			Date (YYYY-MM-DD)	2017-12-08
First Name	Christopher	Last Name	Horgan	Registration Number	40394
Additional Signature may be generated within this form by selecting the Add button.					

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Electronic Patent Application Fee Transmittal

Application Number:	15836824
Filing Date:	
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Filer:	Christopher James Horgan
Attorney Docket Number:	TAAG001CONT1

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
UTILITY FILING FEE (ELECTRONIC FILING)	4011	1	70	70
UTILITY SEARCH FEE	2111	1	300	300
UTILITY EXAMINATION FEE	2311	1	360	360
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**730**

Electronic Acknowledgement Receipt

EFS ID:	31179304
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	09-DEC-2017
Filing Date:	
Time Stamp:	00:39:55
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Fee Worksheet (SB06)	fee-info.pdf	35295 594ac124555e56f507b012c3655444c83bc2bd0f	no	2

Warnings:

Information:	
Total Files Size (in bytes):	35295

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Electronic Patent Application Fee Transmittal

Application Number:	15836824
Filing Date:	
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Filer:	Christopher James Horgan
Attorney Docket Number:	TAAG001CONT1

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
UTILITY FILING FEE (ELECTRONIC FILING)	4011	1	70	70
UTILITY SEARCH FEE	2111	1	300	300
UTILITY EXAMINATION FEE	2311	1	360	360
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**730**

Electronic Acknowledgement Receipt

EFS ID:	31179694
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	09-DEC-2017
Filing Date:	
Time Stamp:	19:01:43
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$730
RAM confirmation Number	121117INTEFSW19024500
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Fee Worksheet (SB06)	fee-info.pdf	35295 72a3896bb8dee75382af6375d2a416a1786663ec	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	35295		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.
National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.
New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
15/836,824	12/09/2017	2425	730	TAAG001CONT1	18	3

CONFIRMATION NO. 9801

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

FILING RECEIPT



Date Mailed: 01/11/2018

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Pouya Taaghol, San Jose, CA;
Vivek Ramanna, San Jose, CA;

Applicant(s)

Pouya Taaghol, San Jose, CA;
Vivek Ramanna, San Jose, CA;

Power of Attorney: None

Domestic Priority data as claimed by applicant

This application is a CON of 14/295,331 06/04/2014 PAT 9843624
which claims benefit of 61/834,807 06/13/2013

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access Application via Priority Document Exchange: No

Permission to Access Search Results: No

Applicant may provide or rescind an authorization for access using Form PTO/SB/39 or Form PTO/SB/69 as appropriate.

If Required, Foreign Filing License Granted: 01/10/2018

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 15/836,824**

Projected Publication Date: Request for Non-Publication Acknowledged

Non-Publication Request: Yes

Early Publication Request: No
** **SMALL ENTITY** **
Title

DISTRIBUTED SOFTWARE DEFINED NETWORKING

Preliminary Class

725

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0175

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 15/836,824

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	70		N/A	
SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	300		N/A	
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	360		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	18 minus 20 =	*	x 40 =	0.00	OR		
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =	*	x 210 =	0.00			
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			0.00			
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				0.00			
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	730		TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT A	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE($)	ADDITIONAL FEE($)	OR	RATE($)	ADDITIONAL FEE($)
AMENDMENT B	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1

CONFIRMATION NO. 9801

IMPROPER CPOA LETTER

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

OC000000096617922

Date Mailed: 01/11/2018

NOTICE REGARDING POWER OF ATTORNEY

This is in response to the power of attorney filed 12/09/2017. The power of attorney in this application is not accepted for the reason(s) listed below:

- The signature(s) of Vivek Ramanna, co-inventor(s) in this application, has/have been omitted. The power of attorney will be entered upon receipt of confirmation signed by said co-inventor(s).

/elbanaybanay/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

112449 7590 06/25/2019

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

EXAMINER	
KRISHNAN, RAJESHWARI	

ART UNIT	PAPER NUMBER
2441	

MAIL DATE	DELIVERY MODE
06/25/2019	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 15/836,824	Applicant(s) Taaghol et al.	
	Examiner RAJESHWARI KRISHNAN	Art Unit 2441	AIA (FITF) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☑ Responsive to communication(s) filed on 12/09/2017..

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☐ This action is **FINAL.** 2b) ☑ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5) ☑ Claim(s) 1-18 is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6) ☐ Claim(s) _____ is/are allowed.

7) ☑ Claim(s) 1-18 is/are rejected.

8) ☐ Claim(s) _____ is/are objected to.

9) ☐ Claim(s) _____ are subject to restriction and/or election requirement

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to **PPHfeedback@uspto.gov.**

Application Papers

10)☑ The specification is objected to by the Examiner.

11)☑ The drawing(s) filed on 12/09/2017. is/are: a)☑ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☑ Notice of References Cited (PTO-892)

2) ☑ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date 12/09/2017.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date _____.

4) ☐ Other: _____.

DETAILED ACTION

Notice of Pre-AIA or AIA Status

1. The present application, filed on or after March 16, 2013, is being examined under the first inventor to file provisions of the AIA.

2. This Office action is in response to US Patent Application 15/836824 filed 12/09/2017.

Priority

3. The Examiner acknowledges Applicant's claim to priority from US Patent Application 14/295331 filed 6/4/2014 and Issued as US Patent US 9843624 on 12/12/2017, which in turn claims priority from US provisional patent application 61/834807 filed on June 13, 2013.

Information Disclosure Statement

4. The information disclosure statement (IDS) submitted on 12/9/2017 is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Reference US 2014/0062976 cited on page 2 of this IDS indicates "Name of the Applicant of Cited Document" to be "Marino, Joseph". But, US 2014/0062976 is by Park et al. This document by Park et al has been considered by the examiner

Specification

5. **The disclosure is objected** to because of the following informalities. Appropriate correction is required.

Table on page 5: "FleXible Manager that manager the …" should be replaced with *FleXible Manager that ~~manager~~ manages the …*

Paragraph [0068] on page 28: "fxSDK simplifies the usage system created APIs" should be replaced with *fxSDK simplifies the usage of system created APIs*

Paragraph [0084] on page 34: "Virtual Fabric (fxVF) provides an abstraction layer for application to communicate with each other" should be replaced with *Virtual Fabric (fxVF) provides an abstraction layer for ~~application~~ applications to communicate with each other*

Paragraph [0088] on page 36: "dRS agents manage the access to the platform resolutions" should be replaced with *dRS agents manage the access to the platform ~~resolutions~~ resources*

Paragraph [0096] on page 39: "verify the authenticity and integrity loaded software" should be replaced with *verify the authenticity and integrity of loaded software*

Claim Objections

6. **Claims 1, 15, 13, 17 objected** to because of the following informalities. Appropriate correction is required.

Regarding claims 1, 15, 17,

The semicolon (";") at the end of the limitation "a programmable cloud device hosting …" should be replaced with a comma (","), because the limitation is followed by a "wherein" clause

the clause "the programmable network device and the programmable cloud device are each powered by" should begin with a "wherein" because this clause recites a further feature of the previously recited *programmable network device* and *programmable cloud device*.

And the same clause "the programmable network device and the programmable cloud device are each powered by" should end with "; and" because it is the penultimate *wherein* clause of the three wherein clauses.

Regarding claim 13, in light of Specification paragraphs [0090-0091], it appears that "load controller is adapted to monitor loads on at least one of the plurality of first network applications and at least one of the plurality of second network applications" should be replaced with *load controller is adapted to monitor loads on the programmable network device and the programmable cloud device ~~at least one of the plurality of first network applications and at least one of the plurality of second network applications~~*.

Regarding claim 13, it appears that "effect change in accordance with thresholds" should be replaced with *~~effect~~ affect change in accordance with thresholds*, because effect is a noun, whereas the verb "affect' would be more appropriate in this limitation.

Claim Interpretation – Contingent clauses not evaluated

7. **In claims 1. 15, 17,** regarding the limitation "wherein at least one of the plurality of first network applications in the programmable network device and at least one of the plurality of second network applications in the programmable cloud device are in secure communication with each other to form a distributed application"

This limitation has not been given patentable weight, and will be ignored while evaluating the claim, because the *devices that host applications* define the scope of the claimed system, but the features of the hosted applications as recited in the *wherein* clause are outside the scope of the system. See MPEP 2111.04.

Note that in order to be given patentable weight, this clause may be amended to recite "how" the devices are configured to enable the applications to have secure communication and to form a distributed application. For instance, "… device further comprise a virtual fabric which provides a secure communication layer" in claim 6 further defines the scope of the claim.

8. **In claims 1, 15, 17,** regarding the limitation, "a sandboxing operating system which facilitates deployment of the plurality of first and second network applications and facilitates upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device"

This limitation has not been given patentable weight, and will be ignored while evaluating the claim. While *devices powered by a sandboxing OS* further defines the

scope of the claimed system, the features of the OS with respect to the deployment/upgrade of applications are outside the scope of the claim. Further, "upgrading applications with substantially no interruption to operation of the devices" is an intended result clause that has no patentable weight in claim interpretation. See MPEP 2111.04.

Note that the limitation may be amended to recite "how" the OS is employed by the devices in the system to upgrade applications, thus making it an active limitation for claim evaluation.

9. **In claims 1, 15, 17,** regarding the limitation, "an application repository storing distributed applications which have been tested for installation in the programmable network device and programmable cloud device"

This limitation has not been given patentable weight, and will be ignored while evaluating the claim. While *application repository storing distributed applications* further defines the scope of the claimed system, the property of the applications such as *being tested for installation in the programmable network device and programmable cloud device* as recited, is outside the scope of the repository and the claim. See MPEP 2111.04. Perhaps, the system includes another component that is configured to provide tested applications for storage in the repository.

10. **In claims 3 and 4,** the limitation, "wherein the at least one of the first/second network applications is responsible for a plurality of functional capabilities of the programmable network/cloud device and is capable of at least one of the group

consisting of …" cannot be given patentable weight because it recites the properties of the applications, which are outside the scope of the claimed *devices which host the applications*, and therefore outside the scope of the claimed system comprising the devices.

11. **In claim 11**, regarding "…device further comprise a distributed content service wherein contents generated by the at least one of the plurality of first network applications may be shared with the at least one of a plurality of second network applications – while "distributed content service" further defines the claimed system, the *wherein* clause cannot be given patentable weight because it defines features of *contents generated by … applications* which are not relevant to the scope of the claimed system. Perhaps the limitation can be amended to recite "how" the *distributed content service* is configured to provide this capability.

12. **In claim 12**, regarding "… device further comprise a distributed notification service wherein notice of a particular event is sent from at least one of the plurality of first network applications to at least one of the plurality of second network applications – while "distributed notification service" further defines the claimed system, the *wherein* clause cannot be given patentable weight because it defines features of *notice of events sent from/to applications* which are not relevant to the scope of the claimed system. Perhaps the limitation can be amended to recite "how" the *distributed notification service* is configured to provide this capability.

Claim Interpretation -112(f)

13. The following is a quotation of 35 U.S.C. 112(f):

> (f) Element in Claim for a Combination. – An element in a claim for a combination may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof, and such claim shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof.

The following is a quotation of pre-AIA 35 U.S.C. 112, sixth paragraph:

> An element in a claim for a combination may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof, and such claim shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof.

14. The claims in this application are given their broadest reasonable interpretation using the plain meaning of the claim language in light of the specification as it would be understood by one of ordinary skill in the art. The broadest reasonable interpretation of a claim element (also commonly referred to as a claim limitation) is limited by the description in the specification when 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, is invoked.

As explained in MPEP § 2181, subsection I, claim limitations that meet the following three-prong test will be interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph:

(A) the claim limitation uses the term “means” or “step” or a term used as a substitute for “means” that is a generic placeholder (also called a nonce term or a non-structural term having no specific structural meaning) for performing the claimed function;

(B) the term “means” or “step” or the generic placeholder is modified by functional language, typically, but not always linked by the transition word “for” (e.g., “means for”) or another linking word or phrase, such as “configured to” or “so that”; and

(C) the term “means” or “step” or the generic placeholder is not modified by sufficient structure, material, or acts for performing the claimed function.

Use of the word “means” (or “step”) in a claim with functional language creates a rebuttable presumption that the claim limitation is to be treated in accordance with 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph. The presumption that the claim limitation is interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, is rebutted when

the claim limitation recites sufficient structure, material, or acts to entirely perform the recited function.

Absence of the word "means" (or "step") in a claim creates a rebuttable presumption that the claim limitation is not to be treated in accordance with 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph. The presumption that the claim limitation is not interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, is rebutted when the claim limitation recites function without reciting sufficient structure, material or acts to entirely perform the recited function.

Claim limitations in this application that use the word "means" (or "step") are being interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, except as otherwise indicated in an Office action. Conversely, claim limitations in this application that do not use the word "means" (or "step") are not being interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, except as otherwise indicated in an Office action.

15. This application includes one or more claim limitations that do not use the word "means," but are nonetheless being interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, because the claim limitation(s) uses a generic placeholder that is coupled with functional language without reciting sufficient structure to perform the recited function and the generic placeholder is not preceded by a structural modifier. Such claim limitation(s) is/are:

Claim 6: a virtual fabric which *provides a secure communication layer* for at the least one of the plurality of first network applications and the at least one of the plurality of second network applications.

Claim 8: a distributed resource service (DRS) which *controls access to a plurality of resources* by the plurality of first network applications and the plurality of the second network applications upon receiving instructions from the application management portal.

Claim 10: a distributed resource service (DRS) which is *capable of providing at least one service* from a group consisting of: exposing application programming interfaces (APIs) to other applications; configuring and managing platform resources; policy enforcement and authorization of applications access to platform resource and other application's APIs; and policy conflict resolution.

Claim 13: load controller is *adapted to monitor loads* on at least one of the plurality of applications and *effect change in accordance with thresholds* received from the application management portal.

If applicant does not intend to have this/these limitation(s) interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, applicant may: (1) amend the claim limitation(s) to avoid it/them being interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph (e.g., by reciting sufficient structure to perform the claimed function); or (2) present a sufficient showing that the claim limitation(s) recite(s) sufficient structure to perform the claimed function so as to avoid it/them being interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph.

16. Because this/these claim limitation(s) is/are being interpreted under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph, they are being interpreted to cover the corresponding structure described in the specification as performing the claimed function, and equivalents thereof.

The Specification discloses that various processes of the dSDN may be implemented by appropriately programmed general purpose or special purpose computing devices, wherein processors execute instructions stored on memory devices, to perform the processes [0031]. Therefore, the specification will be evaluated for disclosure of algorithms/processes/steps that perform the claimed functions.

Regarding claim 13, the specification states that the Load Controller in the dRS could be used to monitor the load on the existing VMs [0091]; a load monitoring application could be notified when the CPU load on a particular fxDevice exceeds a certain threshold the load monitoring application could make smart decisions on reducing the load on the CPU of a fxDevice by forcing handoffs of users to neighboring cells fxDevices [0090].

Therefore, Specification discloses "how" the claimed load controller of claim 13 performs the *monitor loads* and *effect change in accordance with thresholds* functions.

Regarding claims 6, 8, 10, the Specification fails to disclose corresponding algorithms or steps to perform the claimed functions of the virtual fabric and distributed resource service, resulting in rejection under 112(a) and 112(b) as set forth below.

Claim Rejections - 35 USC § 112(a), 112(f)

17. The following is a quotation of the first paragraph of 35 U.S.C. 112(a):

> (a) IN GENERAL.—The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same, and shall set forth the best mode contemplated by the inventor or joint inventor of carrying out the invention.

The following is a quotation of the first paragraph of pre-AIA 35 U.S.C. 112:

> The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same, and shall set forth the best mode contemplated by the inventor of carrying out his invention.

18. **Claims 6, 8, 10 rejected** under 35 U.S.C. 112(a) or 35 U.S.C. 112 (pre-AIA), first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor or a joint inventor, or for pre-AIA the inventor(s), at the time the application was filed, had possession of the claimed invention.

As discussed above, in the "Claim Interpretation -112(f)" section, **"virtual fabric" in claim 6, "distributed resource service (dRS) in claims 8 and 10**, invoke 112(f) interpretation.

The Specification discloses that various processes of the dSDN may be implemented by appropriately programmed general purpose or special purpose computing devices, with processors executing instructions stored on memory devices, thereby performing the processes [0031].

However, the Specification fails to disclose an algorithm to perform each of the claimed specific functions in claims 6, 8, 10. Therefore, these claims lack written description because claims define the invention in functional language, but the

specification does not sufficiently describe how the function is performed or the result is achieved. See MPEP 2161.01(I)

When examining computer-implemented functional claims, examiners should determine whether the specification discloses the computer and the algorithm (e.g., the necessary steps and/or flowcharts) that perform the claimed function in sufficient detail such that one of ordinary skill in the art can reasonably conclude that the inventor possessed the claimed subject matter at the time of filing. It is not enough that one skilled in the art could write a program to achieve the claimed function because the specification must explain "how" the inventor intends to achieve the claimed function to satisfy the written description requirement. See MPEP 2163(II)(A).

Regarding claim 6, the specification states that "Virtual Fabric (fxVF) provides an abstraction layer that enables transparent switching of application and system messages of a dApp between the fxCloud and the fxDevice and between different dApps, by using policies and firewall to enforce security and policies" [0060, 0084], but provides no details about *how* the virtual fabric enforces the policies to provide a secure communication layer or how it enables transparent switching of messages.

Regarding claims 8 and 10, the specification states that the distributed Resources Service (dRS) comprises software agents residing in the fxDevice and the fxCloud devices, which manage access to platform resources and inter-app APIs according to policies [0088], but provides no details about *how* the dRS enforces the policies or what steps the dRS uses to manage access to resources.

Claim Rejections - 35 USC § 112(b), 112(f)

The following is a quotation of 35 U.S.C. 112(b):

> (b) CONCLUSION.—The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the inventor or a joint inventor regards as the invention.

The following is a quotation of 35 U.S.C. 112 (pre-AIA), second paragraph:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

19. **Limitations in claims 6, 8, 10 as listed above invoke 35 U.S.C. 112(f)** or pre-AIA 35 U.S.C. 112, sixth paragraph. However, the written description fails to disclose the corresponding structure, material, or acts for performing the entire claimed function and to clearly link the structure, material, or acts to the function. Therefore, the claims are indefinite and rejected under 35 U.S.C. 112(b) or pre-AIA 35 U.S.C. 112, second paragraph.

As discussed above, in the "Claim Interpretation -112(f)" section, **"virtual fabric" in claim 6, "distributed resource service (dRS) in claims 8 and 10**, invoke 112(f) interpretation.

The Specification discloses that various processes of the dSDN may be implemented by appropriately programmed general purpose or special purpose computing devices, with processors executing instructions stored on memory devices, thereby performing the processes [0031].

However, the Specification fails to disclose corresponding algorithms to perform each of the claimed specific functions in claims 6, 8, 10. Therefore, the claims are indefinite.

Regarding claim 6, the specification states that "Virtual Fabric (fxVF) provides an abstraction layer that enables transparent switching of application and system messages of a dApp between the fxCloud and the fxDevice and between different dApps, by using policies and firewall to enforce security and policies" [0060, 0084], but provides no details about "how" the virtual fabric enforces the policies or how it enables transparent switching of messages.

Regarding claims 8 and 10, the specification states that the distributed Resources Service (dRS) comprises software agents residing in the fxDevice and the fxCloud devices, which manage access to platform resources and inter-app APIs according to policies [0088], but provides no details about how the dRS enforces the policies or what algorithm the dRS uses to manage access.

Applicant may:

(a) Amend the claim so that the claim limitation will no longer be interpreted as a limitation under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph;

(b) Amend the written description of the specification such that it expressly recites what structure, material, or acts perform the entire claimed function, without introducing any new matter (35 U.S.C. 132(a)); or

(c) Amend the written description of the specification such that it clearly links the structure, material, or acts disclosed therein to the function recited in the claim, without introducing any new matter (35 U.S.C. 132(a)).

If applicant is of the opinion that the written description of the specification already implicitly or inherently discloses the corresponding structure, material, or acts and clearly links them to the function so that one of ordinary skill in the art would recognize what structure, material, or acts perform the claimed function, applicant should clarify the record by either:

(a) Amending the written description of the specification such that it expressly recites the corresponding structure, material, or acts for performing the claimed function and clearly links or associates the structure, material, or acts to the claimed function, without introducing any new matter (35 U.S.C. 132(a)); or

(b) Stating on the record what the corresponding structure, material, or acts, which are implicitly or inherently set forth in the written description of the specification, perform the claimed function. For more information, see 37 CFR 1.75(d) and MPEP §§ 608.01(o) and 2181.

Claim Rejections - 35 USC § 112(a)

20. **Claims 1-17 rejected** under 35 U.S.C. 112(a) or 35 U.S.C. 112 (pre-AIA), first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor or a joint inventor, or for pre-AIA the inventor(s), at the time the application was filed, had possession of the claimed invention.

According to the MPEP, original claims may lack written description when the claims define the invention in functional language specifying a desired result but the specification does not sufficiently describe how the function is performed or the result is achieved. See MPEP 2161.01(I). Disclosure of a function or result alone is little more than a wish for possession, it does not satisfy the written description requirement. See MPEP 2163(II)(A)(3)(a).

In independent claims 1, 15, 17, regarding the limitation “sandboxing operating system *which facilitates deployment* of the plurality of first and second network applications *and facilitates upgrades of the first and second network applications with substantially no*

interruption to operation of the programmable network device and programmable cloud device"

The limitation recites desired functions of a sandboxed operating system (fxOS). The Specification merely repeats the desired functions in [0054], "the fxOS is configured to the following … Ninth, it allows hot upgrade of the software and applications on the platform with little to no interruption to the operational aspect of the platform and its applications". But the disclosure fails to sufficiently identify how the fxOS performs the function of hot upgrade with substantially no interruption.

In independent claims 1, 15, 17, regarding the limitation "the programmable network device *verifies authenticity of the upgrades* to the plurality of first network applications and the programmable cloud device *verifies authenticity* of the upgrades to the plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications"

It recites a desired function of verifying authenticity of applications by the claimed devices, but the disclosure fails to sufficiently identify how the function is performed by the devices.

Instead, Specification in [0098] discloses that fxManager (aka, application management portal) uses a hashing algorithm to calculate signature of an application, and verifies authenticity of the application. The specification does not disclose that the *verify authenticity* function is performed by the claimed devices. That is, the specification does not satisfy the written description requirement for the limitation as claimed.

In independent claims 1, 15, 17, regarding the limitation "an application repository storing *distributed applications which have been tested for installation* in the programmable network device and programmable cloud device"

The limitation appears to recite a desired function of storing applications which are tested for installation in devices. The Specification merely repeats the desired functionality in [0053], "using fxManager, the admin may discover new fxApps from the fxStore (aka repository), which presents all the tested and certified vApps as well as showing the supported fxOS version, support hardware platforms, etc."

The disclosure fails to sufficiently identify how the system tests distributed applications for installation. Instead, the Specification recites in [0075], "Prior to installation in the live network, the authorized network admin has probably performed rigorous testing of the application in the test network". That is, the applications are tested by an admin user, and it is improper for the claim language to suggest that the applications have been tested by the claimed system.

In dependent claim 11, regarding the limitation "a distributed content service wherein contents generated by the at least one of the plurality of first network applications may be shared with the at least one of a plurality of second network applications":

The limitation appears to recite a desired function of the distributed content service, such as, to facilitate sharing of contents generated by one of the first network applications with one of the second network applications.

The Specification merely repeats the desired functionality in [0089], "the distributed Content Service (dCS) allows the developers to seamlessly store and share the contents generated by one application with other applications and its associated application in the cloud; dCS simplifies access to the data and brings in storage virtualization to the applications", but fails to sufficiently identify how the Applicant intends to achieve the functionality.

In dependent claim 12, regarding the limitation "a distributed notification service wherein notice of a particular event is sent from at least one of the plurality of first network applications to at least one of the plurality of second network applications":

The limitation appears to recite a desired function of the distributed notification service to facilitate sending of events from one of the first network applications to one of the second network applications.

The Specification merely repeats the desired functionality, "Distributed Notification Service (dNS) by which the applications could be notified of an event to wakeup and/or respond to the event" [0063], "the distributed Notification Service (dNS) is another tool for developers that could wake or ping an application when a particular event has occurred" [0090], but fails to sufficiently identify how the Applicant intends to achieve the functionality.

Claim Rejections - 35 USC § 112(b)

21. **Claim 9 rejected** under 35 U.S.C. 112(b) or 35 U.S.C. 112 (pre-AIA), second paragraph, as being indefinite for failing to particularly point out and distinctly claim the

subject matter which the inventor or a joint inventor, or for pre-AIA the applicant regards as the invention.

Regarding claim 9, "the plurality of resources" in line 1 lacks antecedent basis, because this claim as recited depends from claim 1, whereas "plurality of resources" is recited in claim 8.

Claim Rejections - 35 USC § 103

22. In the event the determination of the status of the application as subject to AIA 35 U.S.C. 102 and 103 (or as subject to pre-AIA 35 U.S.C. 102 and 103) is incorrect, any correction of the statutory basis for the rejection will not be considered a new ground of rejection if the prior art relied upon, and the rationale supporting the rejection, would be the same under either status.

23. This application currently names joint inventors. In considering patentability of the claims the examiner presumes that the subject matter of the various claims was commonly owned as of the effective filing date of the claimed invention(s) absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and effective filing dates of each claim that was not commonly owned as of the effective filing date of the later invention in order for the examiner to consider the applicability of 35 U.S.C. 102(b)(2)(C) for any potential 35 U.S.C. 102(a)(2) prior art against the later invention.

The following is a quotation of 35 U.S.C. 103 which forms the basis for all obviousness rejections set forth in this Office action:

> A patent for a claimed invention may not be obtained, notwithstanding that the claimed invention is not identically disclosed as set forth in section 102, if the differences between the claimed invention and the prior art are such that the claimed invention as a whole would have been obvious before the effective filing date of the claimed invention to a person having ordinary skill in the art to which the claimed invention pertains. Patentability shall not be negated by the manner in which the invention was made.

24. **Claims 1, 2, 5-17 is/are rejected** under 35 U.S.C. 103 as being unpatentable over **Pentikousis** et al ("MobileFlow: Toward Software-Defined Mobile Networks" published in IEEE Communications Magazine • July 2013, hereinafter Pentikousis) in view of **Narayanan** et al (US 2013/0223442 hereinafter Narayanan) and **Shatzkamer** (US 2014/0317293).

Regarding claim 1, **Pentikousis** discloses a system for processing data packets in a network *(MobileFlow: Software Defined Mobile Network (SMDN) architecture, a blueprint for carrier-grade flow-based forwarding [Title, page 1, col. 2 last para])*, the system comprising:

a programmable network device hosting a plurality of first network applications *(MobileFlow forwarding engine MFEE [page 3 col. 2 para3, Fig. 2]; MFEEs are more complex that an OpenFlow switch but much simpler than a router, support virtualized network functions such as DPI [page 4 col. 1 para 2-3]; MFEEs must support carrier-grade functionality such as L3 tunneling and flexible charging [page 4 col. 2 para 1] (i.e. MFEE is similar to a network device such as a router or a switch, and it is programmable depending on the virtualized network functions it hosts))*;

a programmable cloud device hosting a plurality of second network applications *(MobileFlow controller MFC [page 3 col. 2 para3, Fig. 2]; centralized MFCs can steer forwarding in the network [page 4 col. 1 para 3]; mobile network applications running on top of an MFC can function as an MME or the gateway control [page 4 col. 2 para 2] (i.e. the centralized MFC is similar to a cloud device, and it is programmable to host different network functions such as MME or Gateway GW); network function*

virtualization approach: a MobileFlow application implements the control part functionality of each of the EPS elements (e.g., eNodeB-C, SGW-C, PGW-C, S/PGW-C, and GGSN-C). Then each of these applications, which can run in a cloud computing environment interacts with the MFC to manage a respective MFFE [page 6 col1 para 2);

wherein at least one of the plurality of first network applications in the programmable network device and at least one of the plurality of second network applications in the programmable cloud device are in secure communication with each other to form a distributed application *(Each MFFE communicates with an MFC through a lightweight protocol that implements the MobileFlow control interface (Smf) [page 5 col. 1 para 1]; MFFE can be based on a stripped-down mobile gateway that can handle network-layer (L3) tunnel processing in the data plane with all other tunnel control processes moved up to the MFC (i.e. distributed network functions) [page 5 col. 1 last para]; the wireless access MFFE with its eNB-C acts as an eNodeB. The MFFE with its associated SGW-C acts as the SGW, and the MFFE with its associated PGW-C acts as the PGW (i.e. distributed applications on network device and cloud devices) [page 6 col.2 para 3])*;

But, **Pentikousis** does not explicitly disclose, the programmable network device and programmable cloud device are each powered by a sandboxing operating system which facilitates deployment of the plurality of first and second network applications and facilitates upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device.

However, N**arayanan** discloses a similar *software-defined network (SDN) controller 220 and a flow-based switching device 210 [0026, Fig. 2, Fig. 7], wherein the flow-based routing module of the switching device provides a secure sandbox mechanism for separating execution of flow-based routing operations from the rest of the switching device and includes dedicated memory resources and buffering capabilities to facilitate the flow-based network routing [0043].* That is, Narayanan teaches that network devices can employ sandboxing techniques (aka sandboxing OS) to execute applications in a dedicated and protected environment [0043]. It would be easily apparent that cloud devices such as SDN controllers can also employ sandboxing techniques to reap the same benefits.

The rest of this limitation ("a sandboxing operating system which facilitates deployment of applications with substantially no interruption to operation …") cannot be given patentable weight with respect to the claimed system, as explained above in the *Claim Interpretation* sections. However, the Examiner notes that in-service hot-upgrade of software applications is a known benefit of Network Function Virtualization (NFV), as explained in Office actions of the parent application 14/295331 which are inherently incorporated here.

Therefore, it would have been obvious to a person having ordinary skill in the art, before the effective filing date of this application, to modify **Pentikousis** to include sandboxing OS in network and cloud devices, in order to execute applications in a dedicated and protected environment [Narayanan 0043].

But, **Pentikousis** in view of **Narayanan** (hereinafter **Pentikousis-Narayanan**) does not explicitly disclose,

wherein the programmable network device verifies authenticity of the upgrades to the plurality of first network applications and the programmable cloud device verifies authenticity of the upgrades to the plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

an application repository storing distributed applications which have been tested for installation in the programmable network device and programmable cloud device; and

an application management portal coupled to the programmable network device, programmable cloud device and the application repository, wherein the application management portal manages usage of the distributed applications on the programmable network device and programmable cloud device.

However, analogous art by **Shatzkamer** discloses *app store portal for deploying virtualized network functions [Title]; support Network Function Virtualization (NFV) [0005] by providing a digital distribution platform or an "app store" (i.e. repository) for deployment of virtual network functions (VNFs) [0019], vendors or third-party developers communicate with the app store to submit VNFs according to prescribed certification requirements (i.e. tested application, and verify authenticity) [0020]; provide a standardized web based application portal (i.e. application management portal) that can provide point-and-click deployment of the third-party VNFs [0021]. Provide a*

coordinated execution of virtualized network services by an by an orchestrator [0023]; the orchestrator includes a service provisioning module and a life-cycle management controller [0028]; for example, for a new evolved packet service (EPS), the orchestrator can allocate a virtualized eNodeB, a virtualized MME, a virtualized SGW, a virtualized PGW, etc. (i.e. manage use of distributed network applications) [0031].

Therefore, it would have been obvious to a person of ordinary skill in the art, before the effective filing date of the claimed invention, to modify **Pentikousis-Narayanan** to include an application marketplace and an application management portal as taught by **Shatzkamer**, because a standardized web-based app store portal enables a “point-and-click" deployment of virtual network functions, thereby automating the complex tasks of commissioning, capacity planning and management of virtualized network services [Shatzkamer 0001,0004],

Claim 15 is another system claim with limitations similar to those in system claim 1. Claim 15 has additional limitations, namely:

a programmable network device *processing data flows*; and

a programmable cloud device *having a plurality of virtual machines, wherein at least one of the virtual machines has at least one of a plurality of first network applications*.

Pentikousis discloses *MFFEs use the flow rules received from the MFC to encapsulate/decapsulate user traffic and/or forward packets to the transport network (i.e. network device to process data flows) [page 5 col. 1 para 2]; KVM virtual machines*

with different MFCs and applications (i.e. cloud device having VMs running one of the network applications) [page 7 col. 2 para 2]

Shatzkamer also discloses *virtualized infrastructure service provided by hardware resources [0028], to provide execution of the virtualized network function within one or more physical machines [0016].*

It appears that claim 15 does not have any additional features with respect to novelty and inventive steps. Therefore, claim 15 is obvious over **Pentikousis-Narayanan-Shatzkamer**.

Claim 17 is another system claim with limitations similar to those in system claim 1. Claim 17 has additional limitations, namely:

a programmable cloud device *having a plurality of zones, wherein at least one of the plurality of zones has at least one of a plurality of first network applications.*

Shatzkamer discloses that *execution of the virtualized services can be implemented in a data center at a single location or distributed among different geographic locations (i.e. zones) and interconnected via a public or private data network [0025].*

It appears that claim 15 does not have any additional features with respect to novelty and inventive steps. Therefore, claim 15 is obvious over **Pentikousis-Narayanan-Shatzkamer**.

Regarding claim 2, **Shatzkamer** discloses the system of claim 1, wherein each of the programmable network device and programmable cloud device include an access

network interface unit configured to send and receive communications and a processor with a memory associated with the network interface unit *(a physical machine can be configured to execute any of the VNFs … any one of the physical machines can include a network interface circuit, one or more processor circuits, and a memory circuit [0048])*.

The rationale for combining **Pentikousis-Narayanan** and **Shatzkamer** is the same as that discussed above with respect to claim 1.

Regarding claim 5, **Pentikousis** discloses the system of claim 1, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines from the group consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway P-GW) *(An MFFE may include the key functionality of a wireless access node, such as a radio interface to manage radio bearers, wireless access MFFEs can be used in parallel operation with existing eNodeBs [page 4 col.2 last para]; the mobile network applications running on top of an MFC can function as an MME or the gateway control part of a physical box (indicated as GW-C), and thus can interoperate with other legacy MMEs, SGWs, or PGWs [page 4 col. 2 para 2].*

laims 16 and 18 are corresponding system claims with limitations similar to those in claim 5. Therefore, claims 16 and 18 are rejected under the same rationale as claim 5.

Regarding claim 6, **Pentikousis** discloses the system of claim 1, wherein the programmable network device and programmable cloud device further comprise: a virtual fabric which provides a secure communication layer for the at least one of the first network applications and the at least one of the second network applications *()*.

Regarding claim 7, Pentikousis discloses the system of claim 1, further comprising: a second programmable network device which includes the at least one of the first network applications which is in the first programmable network device *(the MFFE can be fully defined in software, to act as different service enablers (e.g., deep packet inspection [DPI], video caching, and optimization) that can be distributed throughout the mobile network, as shown in Fig. 2 [page 1 3 col. 2 last para])*.

Regarding claim 8, **Shatzkamer** discloses the system of claim 1, wherein the programmable network device and programmable cloud device further comprise: a distributed resource service (DRS) which controls access to a plurality of resources upon receiving instructions from the application management portal *(the orchestrator receives request from the portal to deploy a virtual network service, which may communicate with a cloud controller API for controlling VNS provided by hardware resources [0024, 0028])*.

Regarding claim 9, **Shatzkamer** discloses the system of claim [[1]] 8, wherein the resources include at least one of a group consisting of: firewall; data and statistics; storage; compute; load evaluator; general settings; routing engine; virtual machine master controller; power controller; and wireless engine *(the hardware resources include physical server machines, network switching devices and storage devices*

[0028]; NFV can provide a virtualized load balancer [0029]; a life-cycle management controller for starting and suspending VMs, receiving resource statistics [0028]; example services may include Firewall and NAT [0073]).

Regarding claim 10, **Shatzkamer** discloses the system of claim 1, wherein the programmable network device and programmable cloud device further comprise: a distributed resource service which is capable of providing at least one service from a group consisting of: exposing application programming interfaces (APIs) to other applications; configuring and managing platform resources; policy enforcement and authorization of applications access to platform resource and other application's APIs; and policy conflict resolution *(the orchestrator provides life-cycle management of hardware resources [0028], policy mediation and policy charging [Fig. 7, 0046], enforcement of security policies [0081]; the portal presents a list of authorized VNFs and VNSs [0081]).*

Regarding claim 11, **Shatzkamer** discloses the system of claim 1, wherein the programmable network device and programmable cloud device further comprise: a distributed content service wherein contents generated by the at least one of the plurality of first network applications may be shared with the at least one of a plurality of second network applications *(different types of virtualized services can be provided such as video conference or streaming video, IP multimedia system session, etc. all of which require sharing of content between the applications [0031]).*

Regarding claim 12, **Shatzkamer** discloses the system of claim 1, wherein the programmable network device and programmable cloud device further comprise: a distributed notification service wherein notice of a particular event is sent from the at least one of the plurality of first network applications to the at least one of the plurality of second network applications *(The virtual infrastructure manager also can notify the orchestrator of hardware-based capacity issues in either the virtual infrastructure elements or physical hardware resource elements [0031]; the orchestrator can receive an alert from an active VNF container and coordinate the alert among the interdependent active VNF containers across all the applications [0059])*.

Regarding claim 13, **Shatzkamer** discloses the system of claim 10, wherein the DRS further includes a load controller; and wherein the load controller is adapted to monitor loads on at least one of the plurality of first network applications and at least one of the plurality of second network applications and effect change in accordance with thresholds received from the application management portal *(the orchestrator receives hardware load parameters such as CPU, memory or link utilization [0028]; the orchestrator monitors key performance indicator thresholds [0066]; a Virtual Network Function when distributed over multiple VM can make use of an application level load balancing that shall take account of many of the performance indicators to make full and effective use of the available resources [0062]; the app store portal can specify load balancing policies for distributing traffic across multiple physical/logical boundaries [0087])*. Further, **Li** discloses effecting change in accordance with thresholds *(automatically changing the weight or priority of a queue when traffic counts exceed threshold [page 4, column 2])*.

Regarding claim 14, **Shatzkamer** discloses the system of claim 1, wherein the programmable network device, at least one of the plurality of first network applications, the programmable cloud device, and the at least one of the plurality of second network applications each have unique identities *(the orchestrator can maintain IP addresses, server ID and other attributes during service creation [0054]; identification of virtual network functions [0057]; assign Virtual Machine identifiers [0064])*.

25. **Claim 4 is rejected under 35 U.S.C. 103** as being unpatentable over **Pentikousis-Narayanan-Shatzkamer** as applied to claim1,further in view of **Marino** et al (US 2014/0052976 hereinafter **Marino**).

Regarding claim 4, **Pentikousis-Narayanan-Shatzkamer** does not explicitly disclose the system of claim 1, wherein the at least one of the first network applications is responsible for a plurality of functional capabilities of the programmable network device and is capable of at least one of the group consisting of: updating the firmware of the programmable network device; updating default communication protocols of the programmable network device; initializing programmable network device settings; boot loading the programmable network device; providing security digital certificates; providing operating system "(OS) and kernel upgrades; and upgrading and commissioning agents for upgrades of the at least one of the first network applications and related procedures.

However, **Marino** discloses *systems and methods for remotely updating router firmware [0029]. A wireless router receives a firmware update (i.e. first network*

application) from a remote server, and destructively overwrites router firmware in flash memory (i.e. firmware/OS/kernel update/upgrade) in a chunk-wise manner, and then writes a kernel memory before going live with upgraded firmware.

Therefore, it would have been obvious to a person of ordinary skill in the art, before the effective filing date of the claimed invention, to modify **Pentikousis-Narayanan-Shatzkamer** to include first network applications responsible for updating firmware/OS/kernel upgrade of the flexible network device, as taught by **Marino** in order to facilitate remote upgrade of network elements [Marino, 0038].

26. **Claim 5 is rejected under 35 U.S.C. 103** as being unpatentable over **Li-Pentikousis-Narayanan-Shatzkamer** in view of **Baset** et al (US 2013/0263104 hereinafter **Baset**).

Regarding claim 5, **Pentikousis-Narayanan-Shatzkamer** does not explicitly disclose the system of claim 1, wherein at least one of the second network applications is responsible for a plurality of functional capabilities of the programmable cloud device and is capable of at least one of the group consisting of: upgrading and commissioning agents for upgrades of the at least one of the second network applications and related procedures; hypervisor upgrades; virtual machine and/or isolated execution environment lifecycle management; security digital certificate management; operating system (OS) and kernel upgrades; and default communication protocols.

However, **Baset** discloses *a computer software patch is a piece of software that is designed to fix problems with a computer program. For example, the patch may be used to correct security vulnerabilities, correct a programming bug, improve existing functions, or change software behavior of the computer program. Patches may be applied to a hypervisor, an operating system, middleware, and various other computer software applications [0004].*

Therefore, it would have been obvious to a person of ordinary skill in the art, before the effective filing date of the claimed invention, to modify **Pentikousis-Narayanan-Shatzkamer** to include a software or application for upgrading OS/kernel/hypervisor or to manage security issues of a cloud platform as taught by **Baset** in order to facilitate automated patching/upgrading of a computer system in accordance with patch management policies [Baset, Abstract].

Double Patenting

27. The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory double patenting rejection is appropriate where the conflicting claims are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on nonstatutory double patenting

provided the reference application or patent either is shown to be commonly owned with the examined application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement. See MPEP § 717.02 for applications subject to examination under the first inventor to file provisions of the AIA as explained in MPEP § 2159. See MPEP §§ 706.02(l)(1) - 706.02(l)(3) for applications not subject to examination under the first inventor to file provisions of the AIA. A terminal disclaimer must be signed in compliance with 37 CFR 1.321(b).

The USPTO Internet website contains terminal disclaimer forms which may be used. Please visit www.uspto.gov/patent/patents-forms. The filing date of the application in which the form is filed determines what form (e.g., PTO/SB/25, PTO/SB/26, PTO/AIA/25, or PTO/AIA/26) should be used. A web-based eTerminal Disclaimer may be filled out completely online using web-screens. An eTerminal Disclaimer that meets all requirements is auto-processed and approved immediately upon submission. For more information about eTerminal Disclaimers, refer to www.uspto.gov/patents/process/file/efs/guidance/eTD-info-I.jsp.

28. **Claims 1-18 rejected** on the ground of nonstatutory double patenting as being unpatentable over claims 1-14, 18, and 19 of U.S. Patent No. US 9843624. Although the claims at issue are not identical, they are not patentably distinct from each other. Claims 1-14, 18, 19 of the '624 patent recite almost exactly the same invention, with minor variations in language, as claims 1-18 of the present invention. Therefore, claims 1-18 of the present invention are obvious over claims 1-14, 18, 19 of Applicants '624 patent. The mapping of claims is given below.

Claims of present application are obvious over	**Claims of US 9843624**
Independent Claim 1	Independent Claim 1
Dependent Claim 2	Dependent Claim 2
Dependent Claim 3	Dependent Claim 3
Dependent Claim 4	Dependent Claim 4
Dependent Claim 5	Dependent Claim 5
Dependent Claim 6	Dependent Claim 6
Dependent Claim 7	Dependent Claim 7
Dependent Claim 8	Dependent Claim 8
Dependent Claim 9	Dependent Claim 9
Dependent Claim 10	Dependent Claim 11
Dependent Claim 11	Dependent Claim 12

Claims of present application are obvious over	Claims of US 9843624
Dependent Claim 12	Dependent Claim 13
Dependent Claim 13	Dependent Claim 10
Dependent Claim 14	Dependent Claim 14
Independent Claim 15	Independent Claim 18
Dependent Claim 16	Dependent Claim 5
Independent Claim 17	Independent Claim 19
Dependent Claim 18	Dependent Claim 5

Conclusion

29. Any inquiry concerning this communication or earlier communications from the examiner should be directed to RAJESHWARI KRISHNAN whose telephone number is (571)270-5688. The examiner can normally be reached on M-F 8 am - 5 pm.

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at http://www.uspto.gov/interviewpractice.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Wing Chan can be reached on 1-571-272-7493. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic

Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/R.K./
Examiner, Art Unit 2441

/WING F CHAN/
Supervisory Patent Examiner, Art Unit 2441

Notice of References Cited	Application/Control No. 15/836,824	Applicant(s)/Patent Under Reexamination Taaghol et al.	
	Examiner RAJESHWARI KRISHNAN	Art Unit 2441	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-20130223442-A1	08-2013	Narayanan; Rajesh	H04L45/64	370/389
*	B	US-20140317293-A1	10-2014	SHATZKAMER; Kevin	G06F9/455	709/226
*	C	US-20130263104-A1	10-2013	Baset; Salman A.	G06F8/68	717/168
*	D	US-20140052976-A1	02-2014	Marino; Joseph P.	G06F21/56	713/2
	E					
	F					
	G					
	H					
	I					
	J					
	K					
	L					
	M					

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	"MobileFlow - Toward Software Defined Mobile Networks" by Pentikousis et al published in IEEE Communications Magazine • July 2013 (Year: 2013)
	V	Google scholar search results - distributed software defined cellular network (Year: 2019)
	W	Google Scholar search results - network function virtualization eNodeB SGW PGW MME (Year: 2019)
	X	Google scholar search results - Taaghol Pouya (Year: 2019)

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	Art Unit
	RAJESHWARI KRISHNAN	2441

CPC - Searched*		
Symbol	**Date**	**Examiner**
H04W24/02-04	06/16/2019	RK
G06F8/60-68	06/16/2019	RK
G06F9/4406,45533-4856	06/16/2019	RK
H04L 41/0806-0893, 45/00-64,67/10	06/16/2019	RK

CPC Combination Sets - Searched*		
Symbol	**Date**	**Examiner**

US Classification - Searched*			
Class	**Subclass**	**Date**	**Examiner**

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	**Date**	**Examiner**
Performed Inventor and Assignee searches	06/16/2019	RK
EAST search history attached	06/16/2019	RK
Google scholar search for inventor Taaghol Pouya	06/16/2019	RK
Google scholar search for "distributed software defined cellular network"	06/16/2019	RK
Google scholar search for "network function virtualization eNB SGW PGW MME"	06/16/2019	RK

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**

/R.K./ Examiner, Art Unit 2441	

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0217

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	Art Unit
	RAJESHWARI KRISHNAN	2441

/R.K./ Examiner, Art Unit 2441	

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0218

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	RAJESHWARI KRISHNAN	2441

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS										
☐ Claims renumbered in the same order as presented by applicant					☐ CPA		☐ T.D.		☐ R.1.47	
CLAIM		**DATE**								
Final	Original	06/16/2019								
	1	✓								
	2	✓								
	3	✓								
	4	✓								
	5	✓								
	6	✓								
	7	✓								
	8	✓								
	9	✓								
	10	✓								
	11	✓								
	12	✓								
	13	✓								
	14	✓								
	15	✓								
	16	✓								
	17	✓								
	18	✓								

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0219

A virtual SDN-enabled LTE EPC architecture: A case study for S-/P-gateways **functions**

A Basta, W Kellerer, M Hoffmann… - … for Future **Networks** …, 2013 - ieeexplore.ieee.org
… NFV is transforming the way how operators architect their **networks** towards deploying **network** services … Software Defined **Networking** (SDN) and cloud computing. **Virtualized network functions** might run in an operator's cloud environment and might be configured via OpenFlow …
☆ 99 Cited by 162 Related articles All 2 versions »

VPMN: virtual private mobile **network** towards mobility-as-a-service

A Baliga, X Chen, B Coskun, G de los Reyes… - Proceedings of the …, 2011 - cs.utah.edu
… S. Sen, and O. Spatscheck, "Characterizing radio resource allocation for 3g **networks**," in IMC … Bhanage, R. Daya, I. Seskar, and D. Raychaudhuri, "VNTS: a virtual **network** traffic shaper … fairness in 802:16e slices," in IEEE ICC - Wireless and Mobile **Networking** Symposium, 2010 …
☆ 99 Cited by 34 Related articles All 4 versions »

Toward software-defined cellular **networks**

LE Li, ZM Mao, J Rexford - … on Software Defined **Networking**, 2012 - cs.columbia.edu
… Casado, MJ Freedman, J. Pettit, J. Luo, N. Gude, N. McKeown, and S. Shenker, "Rethinking enterprise **network** control," Trans … on **Networking**, vol … G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus **networks**," SIGCOMM CCR …
☆ 99 Cited by 394 Related articles All 18 versions »

Performance evaluation of a SDN/OpenFlow-based Distributed Mobility Management (DMM) approach in **virtualized** LTE systems

L Vallulina, M Karimzadeh… - 2014 IEEE …, 2014 - ieeexplore.ieee.org
… Globecom 2014 Workshop - Cloud Computing Systems, **Networks**, and Applications … multiple DMM controllers during ongoing sessions and within the same DMM transport **network** requires further … This work has been funded by the EU FP7 Mobile Cloud **Networking** (#318109) …
☆ 99 Cited by 31 Related articles All 9 versions

SDN integration in LTE mobile backhaul **networks**

J Costa-Requena - … Conference on Information **Networking** …, 2014 - ieeexplore.ieee.org
… the suitability of the proposed SDN based flatter architecture with LTE **network** elements running … ICT Labs SDN activity to identify the role of SDN technology in LTE mobile **networks**. REFERENCES [1] Cisco visual **networking** index: Global mobile data trac forecast update, 2012 …
☆ 99 Cited by 59 Related articles All 3 versions

NFV: State of the art, challenges and implementation in next generation mobile **networks** (vepc)

H Hawilo, A Shami, M Mirahmadi, R Asal - arXiv preprint arXiv:1409.4149, 2014 - arxiv.org
… Instead, the segment has direct access to the PGW, which routes it to external **networks** … interfaces on GTP leads to better computing performance and use of **network** resources … in terms of minimizing control signaling traffic to avoid congestion on the **networking** infrastructure, it …
☆ 99 Cited by 315 Related articles All 6 versions »

Traffic jam: Handling the increasing volume of mobile data traffic

T Mahmoodi, S Seetharaman - IEEE Vehicular Technology …, 2014 - ieeexplore.ieee.org
… References [1] LE Li, ZM Mao, and J. Rexford, "Toward software-defined cellular **networks**," in Proc. European Workshop Software Defined **Networking**, Oct. 2012, pp. 7–12. [2] Sandvine. (2011, Aug.). Distributed decisions: Hierarchical **network** policy control …
☆ 99 Cited by 46 Related articles All 5 versions

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0220

Applying SDN/OpenFlow in **Virtualized** LTE to Support Distributed Mobility Management (DMM).

M Karimzadeh, L Vaitulina, G Karagiannis - CLOSER, 2014 - pdfs.semanticscholar.org

... The Mobile Cloud **Networking** (MCN) project 0(EU FP7 MCN, 2013), as one of the ... will probably introduce and deploy SDN/OpenFlow Controllers and switches in their operator's transport **networks** ... in Figure 2, a simple OpenFlow **network** is used as transport **network** above the ...

☆ 99 Cited by 22 Related articles All 11 versions »

Moving the mobile evolved packet core to the cloud

J Kempf, B Johansson, S Pettersson... - ... , **Networking** and ..., 2012 - ieeexplore.ieee.org

... The primary **functions** of the EPC with respect to mobile broadband **networking** are mobility ... terminals can cause lots of additional complexity, especially in mobile **networks** [8]. With ... Meanwhile, other termnals that are not infected can enjoy mobile **network** service without the ...

☆ 99 Cited by 161 Related articles All 4 versions

Cache in the air: Exploiting content caching and delivery techniques for 5G systems

X Wang, M Chen, T Taleb, A Ksentini... - IEEE ..., 2014 - ieeexplore.ieee.org

... The deployment of caches within EPC using conventional content delivery **networking** (CDN) techniques has ... Furthermore, caching at the edge of the **network** has the effect of smoothing traffic ... 1, 3G and 4G cellular **networks**, as well as the expected 5G **networks**, generally have ...

☆ 99 Cited by 698 Related articles All 9 versions

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0221

Mobility protocol switching for wireless networks

P Taaghol - US Patent App. 11/680,790, 2008 - Google Patents

Techniques for mobility protocol switching for wireless networks are described. An apparatus may comprise a convergence gateway having a protocol converter, the protocol converter to receive as input a first set of control plane protocol signals and a first set of user ...

☆ Cited by 23 Related articles All 2 versions

Distributed software defined networking

P Taaghol, V Ramanna - US Patent 9,843,624, 2017 - Google Patents

Abstract The Distributed Software Defined Network (dSDN) disclosed herein is an end-to-end architecture that enables secure and flexible programmability across a network with full lifecycle management of services and infrastructure applications (fxDeviceApp). The dSDN ...

☆ Cited by 24 Related articles All 2 versions

Vertical handover composite quality measures

P Taaghol, V Gupta - US Patent 8,565,744, 2013 - Google Patents

A composite quality measure helps determine a target for a vertical handover between different connection types. Various measurements are normalized, weighted and summed to produce composite quality measures. A composite quality measure may be compared to a ...

☆ Cited by 20 Related articles All 4 versions

Wireless network handover with single radio operation

P Taaghol - US Patent App. 12/126,737, 2009 - Google Patents

Briefly, in accordance with one or more embodiments, a handover of a wireless device from a first network to a second network may be implemented via operation of a single radio in the wireless device. A network architecture capable of implementing such a single radio ...

☆ Cited by 20 Related articles All 2 versions

Method and apparatus for virtualization of wireless network

P Taaghol - US Patent App. 12/858,684, 2012 - Google Patents

A method, apparatus and a system to share a wireless spectrum by mobile devices of different operators. The sharing of the wireless spectrum is done by providing a first virtual base station to communicate with mobile devices of a first operator and providing a second ...

☆ Cited by 17 Related articles All 2 versions

Mobile station localization apparatus, systems, and methods

P Taaghol, M Venkatachalam - US Patent 7,787,891, 2010 - Google Patents

Embodiments herein may use first and second path loss values L1 and L2 to calculate first and second distances R1 and R2 between the MS and first and second base stations BS1 and BS2, respectively, in aa reduction in signal strength of signals as transmitted by the BS1 ...

☆ Cited by 18 Related articles All 5 versions

Emergency call services for wireless network roaming

P Taaghol, M Venkatachalam - US Patent 8,094,651, 2012 - Google Patents

A method and apparatus for emergency call services for wireless network roaming. An embodiment of a method includes receiving a request for an emergency telephone call from a mobile device, the mobile device communication over the Internet, the mobile device ...

☆ Cited by 17 Related articles All 5 versions

Gateway association

PK Jain, M Venkatachalam, P Taaghol - US Patent 8,913,586, 2014 - Google Patents

Embodiments provide components of a communication network that may facilitate efficient optimization of the communication network. In various embodiments, network gateway selection for incoming subscriber stations may be allocated to one or more components that ...

☆ Cited by 17 Related articles All 4 versions

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0222

Generic bootstrapping protocol (gbp)

P Taaghol - US Patent App. 12/941,205, 2011 - Google Patents

Embodiments of methods and apparatus for discovering a bootstrapping server by a mobile station, and bootstrapping the mobile station are disclosed. In some embodiments, a mobile station discovers a remotely disposed bootstrapping server; transmits wirelessly a bootstrap ...

☆ 99 Cited by 15 Related articles All 2 versions »

Deregistration with context retention for inter-RAT operations

..., M Venkatachalam, A Sharaga, P Taaghol - US Patent ..., 2015 - Google Patents

A method to perform deregistration with context retention is presented. In one embodiment, the method includes storing a context associated with a mobile station in response to a deregistration request from the mobile station. The method includes determining whether or ...

☆ 99 Cited by 14 Related articles All 4 versions »

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0223

PTO/SB/08a (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 1 of 3

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	Taaghol
Art Unit	
Examiner Name	
Attorney Docket Number	TAAGCONT1

U. S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number / Number-Kind Code[2] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
		US-8,528,108	09-2013	Ammer, Gerhard	
		US-8,756,595	06-2014	Arcilla, Andre	
		US-9,226,333	12-2015	Steiner, Moritz	
		US-20090300057	12-2009	Friedman, Nathaniel	
		US-20140241158	08-2014	Anthony, Jr., Bruce	
		US-20110075611	03-2011	Choi, Hyung-Nam	
		US-7,630,706	12-2009	Weigand, David	
		US-20130007254	01-2013	Fries, Robert	
		US-20090113327	04-2009	Pogrebinsky, Vladimir	
		US-20080209491	08-2008	Hasek, Charles	
		US-20140317293	10-2014	Shatzkamer, Kevin	
		US-20140254373	09-2014	Varma, Subir	
		US-20140259012	09-2014	Nandiali, Vishwamitra	
		US-20150173121	06-2015	Miklos, Gyorgy	
		US-20070078988	04-2007	Miloushev, Vladimir	
		US-2009027677	11-2009	Nickolov, Peter	
		US-20020016954	02-2002	Charities, Dietrich	
		US-7114021	09-2006	Seshadri	
		US-8132166	03-2012	DeHaan	

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document / Country Code[3]-Number[4]-Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]

Examiner Signature	/RAJESHWARI KRISHNAN/	Date Considered	06/11/2019

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /R.K./

PTO/SB/08a (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 2 of 3

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	Taaghol
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001CONT1

U. S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code[2] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
		US-20040107417	06-2004	Chis	
		US-20050246703	11-2005	Ahonen	
		US-20090007091	01-2009	Appiah	
		US-20040237081	11-2004	Homilier, Daniel	
		US-20120300615	11-2012	Kempf, James	
		US-20130160147	06-2013	Draluk, Vadim	
		US-20130263104	10-2013	Baset, Salman	
		US-20140062976	02-2014	Marino, Joseph	
		US-8,650,299	02-2014	Huang, Alex	
		US-8,843,933	09-2014	Holler, Anne Marie	
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document Country Code[3]-Number[4]-Kind Code[5] (if known)	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]

Examiner Signature	/RAJESHWARI KRISHNAN/	Date Considered	06/11/2019

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /R.K./

PTO/SB/08b (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 3 of 3

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	TAAGHOL
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001CONT1

NON PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
		Bansal, et al., "OpenRadio: A Programmable Wireless Dataplane", published in HotSDN'12, Aug. 13, 2012, Helsinki, Finland	
		Stallings et al., "Software-Defined Networks and Open Flow" published in The Internet Protocol Journal, Vol. 16, No. 1, March 2012	
		Luo et al., "Software Defined Networking across Distributed Datacenters over Cloud" published 2013 IEEE International Conference on Cloud Computing Technology and Science	
		McKeown et al., "OpenFlow:Enabling Innovation in Campus Networks", March 14, 2008.	
		Li et al., "Toward Software-Defined Cellular Networks", March 14, 2008.	
		"Network Functions Visualization", White Paper published Oct. 22-24, 2012, at the "SDN and OpenFlow World Congress", Damstadt-Germany	
		Yap et al., "OpenRoads: Empowering Research in Mobile Networks" published in 2010 ACM SIGCOMM Computer Communication Review	
		"A Short Introduction To Cloud Platforms " by David Chappell, Aug 2008	

Examiner Signature	/RAJESHWARI KRISHNAN/	Date Considered	06/11/2019

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /R.K./

Toward **software-defined cellular networks**

LE Li, ZM Mao, J Rexford - … on **Software Defined Networking**, 2012 - cs.columbia.edu

... SDN can enable carriers to **distribute** data-plane rules over multiple, cheaper ... and extensions to controller platforms, **network** switches, and base stations to enable **software defined cellular networks** ... [8] F. Kuhn and R. Wattenhofer, "On the complexity of **distributed** graph coloring ...

☆ 99 Cited by 394 Related articles All 18 versions »

Toward secure **distributed** spectrum sensing in cognitive radio **networks**

R Chen, JM Park, YT Hou... - IEEE Communications …, 2008 - ieeexplore.ieee.org

... Fur- thermore, future DSA applications may be extended into other licensed bands such as those used by **cellular networks** ... **Networking** Technologies for **Software Defined** Radio **Networks**, Sept ... [8] PK Varshney, **Distributed** Detection and Data Fusion, Springer-Verlag, 1997 ...

☆ 99 Cited by 411 Related articles All 12 versions

Software-defined networking (SDN) and **distributed** denial of service (DDoS) attacks in cloud computing environments: A survey, some research issues, and …

Q Yan, FR Yu, Q Gong, J Li - IEEE communications surveys & …, 2015 - ieeexplore.ieee.org

... Denial of Service (DoS) attacks and **Distributed** Denial of Service (DDoS) flooding attacks are the main methods to destroy ... ONF has provided the most explicit and well received **definition** of SDN as follows: "In the SDN architecture, the ... What Is **Software-Defined Networking** ...

☆ 99 Cited by 349 Related articles

A survey of **software-defined networking**: Past, present, and future of programmable **networks**

BAA Nunes, M Mendonca, XN Nguyen... - … Surveys & Tutorials, 2014 - ieeexplore.ieee.org

... in Palette is to partition end-to-end policies into sub tables and then **distribute** them over ... is not **defined** by OpenFlow, it is necessary for any type of **distribution** or redundancy ... the idea further by attempting to maintain a logically centralized but physically **distributed** control plane ...

☆ 99 Cited by 1538 Related articles All 22 versions

SoftRAN: **Software defined** radio access **network**

A Gudipati, D Perry, LE Li, S Katti - … topics in **software defined networking**, 2013 - dl.acm.org

... Further, **distributed** coordination algorithms tend to become more complex, since often they require ... as a result, 3GPP (the LTE standards body) is working towards **defining** coordi- nation ... being able to rapidly implement such protocols through our **software-defined** architecture ...

☆ 99 Cited by 444 Related articles All 14 versions »

Software-defined wireless **networking**: centralized, **distributed**, or hybrid?

M Abolhasan, J Lipman, W Ni, B Hagelstein - IEEE Network, 2015 - ieeexplore.ieee.org

... To accomplish scalability, we can **distribute** SDN controllers and divide the **network** ... existing SDN with fully centralized control and purely **distributed** routing approaches ... K. Pentikousis, Y. Wang, and W. Hu, "Mobileflow: Toward **Software- Defined Mobile Networks**," IEEE Commun ...

☆ 99 Cited by 44 Related articles All 2 versions

An architecture for **software defined** wireless **networking**

CJ Bernardos, A De La Oliva… - IEEE Wireless …, 2014 - eprints.networks.imdea.org

... SDN) is currently being considered as an alternative to classic **distributed** approach- es ... on the different control and user plane entities, which require **software**/firmware updates (if ... best dynamic use of available resources, considering current load, users' **distribution**, and **network** ...

☆ 99 Cited by 265 Related articles All 13 versions »

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0227

Mobileflow: Toward **software-defined mobile networks**

K Pentikousis, Y Wang, W Hu - IEEE Communications magazine, 2013 - Citeseer

... However, as new technolo- gies such as content **distribution** and cloud com- puting enter the **mobile** operator ... service enablers (eg, deep packet inspection [DPI], video caching, and optimization) that can be **distributed** throughout the ... Figure 2. **Software defined mobile network** ...

☆ 99 Cited by 336 Related articles All 3 versions »

Publish/subscribe-enabled **software defined networking** for efficient and scalable IoT communications

A Hakiri, P Berthou, A Gokhale, S Abdellatif - arXiv preprint arXiv ..., 2017 - arxiv.org

... Internet of Things," Internet of Things (WF-IoT), 2014 IEEE World Forum on , Mar 2014, pp.89-93 [5] Object Management Group, "Data-**Distribution** Service for ... 2013 [8] Y. Li, H. Wang, M. Liu; B. Zhang, H. Mao, "**Software Defined Networking** for **Distributed** Mobility Management ...

☆ 99 Cited by 88 Related articles All 11 versions »

SoftAir: A **software defined networking** architecture for 5G wireless systems

IF Akyildiz, P Wang, SC Lin - Computer Networks, 2015 - Elsevier

... 5, which determines how to **distribute** non-conflicting **network** resource blocks among virtual **network** ... **network** hypervisor, the distance vectors that depend on the user spatial **distribution** of the ... The third enabling tool that will be developed is a **distributed** traffic classifier (Fig ...

☆ 99 Cited by 163 Related articles All 8 versions

Related searches

security in software defined networks

software defined **radio access** network

software defined **wireless** networking

software defined networking **approach**

software defined networking **architecture**

software defined networking **iot communications**

flow installation software defined networks

software defined networking **cloud computing**

networking **sdn** distributed **denial of service**

secure distributed **spectrum sensing radio** networks

things devices software defined networks

wireless mobile networks

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0228

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S308	3	"distributed software defined network"	USPAT	OR	OFF	2019/06/11 08:51
S309	11	"distributed software defined network"	US-PGPUB; USPAT; EPO	OR	OFF	2019/06/11 08:55
S310	3	"distributed software defined network" and (@ad @rlad) <= "20130613"	US-PGPUB; USPAT; EPO	OR	OFF	2019/06/11 09:20
S311	250	sandbox$3 adj2 ("operating system" OS)	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/11 09:20
S312	69	sandbox$3 adj ("operating system" OS)	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/11 09:21
S313	33	sandbox$3 adj ("operating system" OS) and (@ad @rlad) <= "20130613"	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/11 09:21
S316	29	("0000007" | "20020016954" | "20040107417" | "20040237081" | "20050246703" | "20070078988" | "20080209491" | "20090007091" | "20090027677" | "20090113327" | "20090300057" | "20110075611" | "20120300615" | "20130007254" | "20130160147" | "20130263104" | "20140062976" | "20140241158" | "20140254373" | "20140259012" | "20140317293" | "20150173121" | "7114021" | "8132166" | "8528108" | "8650299" | "8756595" | "8843933" | "9226333").PN.	US-PGPUB; USPAT	OR	OFF	2019/06/11 16:01
S318	1	("20140062976").PN.	US-PGPUB; USPAT	OR	OFF	2019/06/11 16:09
S328	21	("10069648" | "20070053334" | "20090228589" | "20120204224" | "20130054761" | "20130094472" | "20130139243" | "20140341199" | "20150124622" | "20150200859" | "20150382240" | "20160085641" | "20160100340" | "20160127889" | "20160156513" | "20160182378" | "20160301595" | "20160315785" |	US-PGPUB; USPAT	OR	OFF	2019/06/13 10:12

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0229

		"20160381146" | "20180248713" | "9264944").PN.				
S329	7	("20070115906" | "20090094611" | "20120106333" | "20130054761" | "20130297798" | "20150055623" | "20150124815").PN.	US-PGPUB; USPAT	OR	OFF	2019/06/13 10:14
S330	10	("20120300615" | "20130242727" | "20130272127" | "20130329601" | "20140038654" | "20140101726" | "20140269284" | "20140348130" | "20150063346" | "8521905").PN.	US-PGPUB; USPAT	OR	OFF	2019/06/13 10:16
S331	0	(S328 or S329 S330) and (distributed adj2 software-defined adj2 network$3)	US-PGPUB; USPAT	OR	OFF	2019/06/13 10:17
S332	18	(S328 or S329 or S330) and ((distributed near5 software-defined adj2 network$3) OR "virtual network function" OR "network functions virtualization" OR open$1flow)	US-PGPUB; USPAT	OR	OFF	2019/06/13 10:19
S333	2	("20110044218" "20160150448").pn.	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/13 10:44
S334	11	13/536838	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/13 13:09
S335	68	("2014/0317293").URPN.	USPAT	OR	OFF	2019/06/13 13:12
S337	148	(((distributed near5 software-defined adj2 network$3) OR "virtual network function" OR "network functions virtualization" OR open$1flow)) and (eNB e$1NodeB S$1GW P$1GW MME) and (@ad @rlad) <= "20130613"	US-PGPUB; USPAT	OR	OFF	2019/06/13 13:18
S338	5	("0000007" | "20020016954" | "20040107417" | "20040237081" | "20050246703" | "20070078988" | "20080209491" | "20090007091" | "20090027677" | "20090113327" | "20090300057" | "20110075611" | "20120300615" | "20130007254" | "20130160147" | "20130263104" | "20140062976" | "20140241158" | "20140254373" | "20140259012" | "20140317293" | "20150173121" | "7114021" | "8132166" | "8528108" | "8650299" | "8756595" | "8843933" | "9226333").PN. and (((distributed near5 software-defined adj2 network$3) OR "virtual network function" OR "network functions virtualization" OR open$1flow)) and (eNB e$1NodeB S$1GW P$1GW MME)	US-PGPUB; USPAT	OR	OFF	2019/06/13 13:19
S339	2	13/912658	US-PGPUB; USPAT	OR	OFF	2019/06/13 13:20
S340	6	("0000007" | "20020016954" | "20040107417" | "20040237081" | "20050246703" | "20070078988" |	US-PGPUB; USPAT	OR	OFF	2019/06/13 13:22

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0230

		"20080209491" | "20090007091" | "20090027677" | "20090113327" | "20090300057" | "20110075611" | "20120300615" | "20130007254" | "20130160147" | "20130263104" | "20140062976" | "20140241158" | "20140254373" | "20140259012" | "20140317293" | "20150173121" | "7114021" | "8132166" | "8528108" | "8650299" | "8756595" | "8843933" | "9226333" "20140362775").PN. and (((distributed near5 software-defined adj2 network$3) OR "virtual network function" OR "network functions virtualization" OR open$1flow)) and (eNB e$1NodeB S$1GW P$1GW MME)				
S341	5	Pentikousis.inv.	US-PGPUB; USPAT	OR	OFF	2019/06/13 14:28
S342	0	Pentikousis.inv. and smdn	US-PGPUB; USPAT	OR	OFF	2019/06/13 14:29
S343	3	Pentikousis.inv. and "software defined"	US-PGPUB; USPAT	OR	OFF	2019/06/13 14:29
S344	79	((Pouya) near2 (Taaghol)).INV.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/06/13 14:38
S345	3	((Vivek) near2 (Ramanna)).INV.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/06/13 14:38
S346	1	S344 and S345	US-PGPUB; USPAT; USOCR	OR	OFF	2019/06/13 14:38
S347	1	MobileFlow and (@ad @rlad) <= "20130613"	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/13 14:56
S348	24	("Extensible Messaging and Presence Protocol" or XMPP or Jabber) and ("software defined network" "software defined mobile network" "software defined cellular network") and (@ad @rlad) <= "20130613"	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/13 15:03
S349	86	("Extensible Messaging and Presence Protocol" or XMPP or Jabber) and ("software defined network" "software defined mobile network" "software defined cellular network")	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/13 15:04
S350	62	(sandbox$3) and ("software defined network" "software defined mobile network" "software defined cellular network")	US-PGPUB; USPAT; USOCR; EPO	OR	OFF	2019/06/13 16:07
S351	5	(sandbox$3) and ("software defined network" "software defined mobile network" "software defined cellular network") and (@ad @rlad) <= "20130613"	US-PGPUB; USPAT; USOCR;	OR	OFF	2019/06/13 16:08

			EPO			
S352	79	((Pouya) near2 (Taaghol)).INV.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/06/16 11:08
S353	3	((Vivek) near2 (Ramanna)).INV.	US-PGPUB; USPAT; USOCR	OR	OFF	2019/06/16 11:08
S354	1	S352 and S353	US-PGPUB; USPAT; USOCR	OR	OFF	2019/06/16 11:08
S355	1	((Pouya) near2 (Taaghol)).INV. and distributed and flexible	USPAT	OR	OFF	2019/06/16 11:08
S356	1	((Pouya) near2 (Taaghol)).INV. and (distributed near9 "software defined" near2 network)	USPAT	OR	OFF	2019/06/16 11:09
S357	1	((Pouya) near2 (Taaghol)).INV. and ("software defined" near2 network)	USPAT	OR	OFF	2019/06/16 11:09

6/16/2019 11:10:34 AM
C:\Users\rkrishnan\Documents\EAST\Workspaces\14295331_15836824 Taaghol Distributed SW Defined Network.wsp

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0232

MobileFlow: Toward Software-Defined Mobile Networks

Kostas Pentikousis, Yan Wang, and Weihua Hu, Huawei Technologies

ABSTRACT

Mobile carrier networks follow an architecture where network elements and their interfaces are defined in detail through standardization, but provide limited ways to develop new network features once deployed. In recent years we have witnessed rapid growth in over-the-top mobile applications and a 10-fold increase in subscriber traffic while ground-breaking network innovation took a back seat. We argue that carrier networks can benefit from advances in computer science and pertinent technology trends by incorporating a new way of thinking in their current toolbox. This article introduces a blueprint for implementing current as well as future network architectures based on a software-defined networking approach. Our architecture enables operators to capitalize on a flow-based forwarding model and fosters a rich environment for innovation inside the mobile network. In this article, we validate this concept in our wireless network research laboratory, demonstrate the programmability and flexibility of the architecture, and provide implementation and experimentation details.

INTRODUCTION

Carrier networks rely on pricy, tightly integrated, and monolithic machinery which is neither easy to configure optimally nor troubleshoot. Carrier equipment follows widely accepted network standards, such as those introduced by the Third Generation Partnership Project (3GPP). However, typical implementations rely on vendor-specific hardware platforms. Operators often end up in vendor lock-ins as entire areas must be covered with equipment from the same vendor for maximum efficiency. These problems are manifested in the unlicensed spectrum wireless industry too. For example, the control and provisioning of wireless access points (CAPWAP) standard (RFC 5415) enables a single access controller to manage numerous WiFi access points, possibly procured from a variety of vendors. In practice, CAPWAP multivendor interoperable deployments are uncommon. More critically, the closed nature of modern network equipment prevents the research community from exploring new paradigms using real-world equipment. For mobile operators, although new revenue opportunities arise, such as machine-type communication (MTC), standardization takes years from concept definition to commercial equipment that can realize the service becoming widely available. In short, the carrier network paradigm we have come to trust has reached its limits.

On these grounds, software-defined networking (SDN) emerged, aiming at a shift toward a flow-centric model that employs inexpensive hardware [1], a logically centralized network controller, and assorted applications that utilize controller-exposed information to orchestrate service delivery in the network [2]. SDN takes cues from the way we run networks today: modern networks, be they campus, data center, enterprise, or carrier networks, follow domain-centered deployment and operation. In this paradigm, each domain could develop its own network services, addressing specific user needs. Presently, this can only be done on an end-to-end basis (i.e., on top of IP), as only vendors can modify the highly sophisticated but primarily hardware-based network elements. Instead, SDN advocates the use of simple but software-programmable network switches. This model can take advantage of modern agile programming methodologies, which enable network software to be developed, enhanced, and upgraded at much shorter cycles than what we experience today with state-of-the-art hardware-centric network machinery.

The first generation of SDN development is closely associated with OpenFlow [2], but the lasting contribution of this early phase work is ushering in radical new thinking in the way we design and realize networks. SDN explicitly separates the control and data planes in a manner more daring than other carrier-grade architectures, such as the 3GPP Evolved Packet Core (EPC). Moreover, SDN has triggered significant interest in network function virtualization (NFV), as indicated by the latest developments in standardization organizations such as the Internet Engineering Task Force (IETF) and European Telecommunications Standards Institute (ETSI). In this context, we advance the state of the art in SDN through three contributions. First, we introduce the software-defined mobile network (SDMN) architecture, a blueprint for carrier-grade flow-based forwarding and a rich environment for innovation at the

0163-6804/13/$25.00 © 2013 IEEE

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0233

core of the mobile network. The first publication of concept validation laboratory results is our second contribution. The hindsight and implementation details provided should be of great interest to the research community at large. Finally, backed by experimental testing, we explore the programmability, flexibility, and openness of the proposed architecture, detailing how mobile networks can be developed and deployed in the future.

The remainder of this article is organized as follows. After reviewing pertinent related work, we provide a brief background on EPC. We then present SDMN, detailing the core components, and explain the different models supported. Finally, we illustrate our research laboratory experiments and conclude the article.

Related Work

Since the introduction of OpenFlow [2], prominent SDN work has addressed several topics such as network virtualization [1, 3], data center and cloud networking [4, 5], acceleration in value-added network service development [6], and network management and control platforms [7], taking an implementation- and experimentation-oriented approach from the beginning. Wireless networking has also been considered, although salient work in this area, to date, has focused primarily on campus-like wireless deployments. The feasibility of mobility management with vertical handovers between IEEE 802.11 and IEEE 802.16 networks was explored in [8]. Although this work is indicative of future developments in wireless networking, it does not address the challenges mobile operators face when extending a model originally designed for wired networks to the wireless domain. Specifically, carrier-grade software-defined mobile networks have not received the attention they deserve given their ubiquity and importance in connecting the majority of Internet users around the world.

Mobile operators are interested in mature limited-risk technologies that are highly optimized to make the most out of scarce (licensed) wireless resources, scalable to handle billions of users as today's networks can, and ready to foster innovation inside the network. A recent position paper on software-defined cellular networks [9] makes the case that SDN can simplify cellular networks and lower management costs. We concur with Li *et al.* [9] that as one introduces SDN in a mobile network, the challenges lying ahead are significant. Unfortunately, the authors do not provide an architectural blueprint detailing their proposal; nor do they report experimental results. On the other hand, Kempf *et al.* [10] present a detailed study on the evolution of EPC toward a model where the control plane can be "lifted up" from the core network elements and henceforth reside in a data center. The authors review the current EPC architecture and document their implementation, which is based on OpenFlow 1.2 protocol extensions, at a sufficient level of detail so that it can be replicated by others in the SDN community. The article, which we view as complementary to our work, emphasizes the role that data centers can play in future carrier networks and lists possible architectural modifications to EPC, but is very terse on prototype details. Indeed, the authors' experience with using the prototype appears as a work in progress, and the article does not provide experimentation results. The work presented in this article can coincide with that described in [10], but takes a different approach that is not as closely knit to the use of data centers while deploying the proposed architecture. NFV is currently expected to be a driving force in future carrier network standardization, and there are many different paths one can follow in both the system design phase and actual deployment. We argue in this article that SDMN is demonstrably flexible enough to accommodate numerous ways forward.

This article advances the state of the art in the public peer-reviewed SDN literature by describing a concrete architectural proposal that applies the core SDN principles to mobile carrier networks. We detail the development and use of a proof-of-concept prototype to validate the key components of the architecture, following the SDN tradition of implementation-oriented experimentation.

Mobile Broadband Today

EPC is a special-purpose, flat, all-IP architecture standardized by 3GPP [11–13]. EPC, which along with the evolved Universal Mobile Telecommunications System (UMTS) terrestrial radio access network (E-UTRAN) forms the foundation of the Evolved Packet System (EPS) and fourth generation (4G) networks, is currently deployed by dozens of operators around the world. EPC made significant steps forward in terms of embracing packet-switched networking, aiming at reduced complexity and increased scalability through data and control plane separation. EPC acts as a unifying routing fabric in the core network used by different 3GPP-standardized radio technologies as well as non-3GPP access technologies such as IEEE 802.11, among others.

EPC is a highly optimized mobile service overlay that capitalizes on IP- or Ethernet-based transport to deliver a diverse range of services through secure communications with quality of service (QoS) guarantees and seamless mobility support [11–13]. Figure 1 illustrates the key entities, functional elements, and interfaces involved in typical EPS operation during a voice call, video streaming, or online access to web content. We do not delve into the details of EPS operation due to space constraints (interested readers should consider [11–13, references therein]), but note that traffic from/to the user equipment (UE) is effectively directed through a set of general packet radio service (GPRS) Tunneling Protocol (GTP)/Proxy Mobile IP (PMIP) tunnels from the attached eNB to the serving gateway (SGW) to the packet data network (PDN) gateway (PGW) and onward to the corresponding PDN; see [11] for more details. As illustrated in Fig. 1, user packets are encapsulated/decapsulated and processed successively by different functional elements associated with operational and management functions such as

EPC is a special-purpose, flat, all-IP architecture standardized by 3GPP. EPC, which along with E-UTRAN, forms the foundation of the EPS and 4G networks, is currently deployed by dozens of operators around the world.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0234

A top-level requirement for SDMN is to provide maximum flexibility, openness, and programmability to future carriers without mandating any changes in UE. In this way, operators can innovate inside their domain without having to depend on either over-the-top (OTT) service providers or UE vendors to support their innovations.

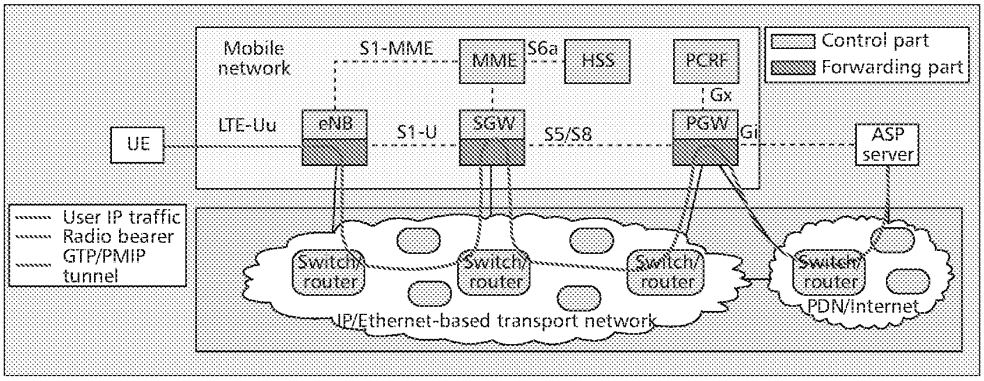


Figure 1. *Evolved packet system.*

charging, policing, authentication, authorization, and accounting (AAA) operations, mobility management, and so on. Each "box" in this architecture typically stands for a physical network device, often powered by specialized software and hardware, adhering to carrier-grade operation requirements. This orientation toward specialized hardware stems, in part, from the evolution of telecommunication networks [14, references therein].

This mobile network architecture has been in general successful as manifested by the number of mobile subscribers and the multibillion dollar industry around it. However, as new technologies such as content distribution and cloud computing enter the mobile operator domain, the complex nature of mobile broadband technologies starts to become an impediment to sustainable future growth. EPC is a fit-for-purpose closed system based on standardized interfaces, where every component performs specific functions, and each of the dozens of interfaces has a unique definition. Further evolution of the mobile broadband architecture along this line of thinking has certain drawbacks. First, operators will have to deal with higher capital and operational expenditures (CAPEX/OPEX) at a time when average revenue per user (ARPU) is decreasing. Progress down this evolutionary path, based on this so-called stovepipe design model, is bound to incur an increasing rate of costs in the future. Meanwhile, operations staff have to maintain many types of network elements and follow the evolution of dozens of interfaces. Second, as higher CAPEX/OPEX forces some operators to refrain from investing further, those who do invest face long time-to-market periods as it is taxing to add new features. Interested operators must push a whole industry to standardize these new features, and then wait for vendors to actually implement them. This process does assure quality standards development, but prevents operators who are willing to endeavor first into new domains from doing so in a fast pace. Finally, a direct consequence of the specialization of each network element, which is defined solely for EPC, is limited openness and flexibility. Once purchased, EPC elements can be controlled through standard interfaces, but cannot be extended through the use of open interfaces or application programming interfaces (APIs). If the operator expects that new kinds of network services would appeal to its subscribers, existing equipment may be of little use in developing and deploying said services.

Software-Defined Mobile Network

In order to address the challenges introduced in the previous sections, and motivated by the core principles of the SDN paradigm, we set out to define what we refer to as the software-defined mobile network (SDMN). A top-level requirement for SDMN is to provide maximum flexibility, openness, and programmability to future carriers without mandating any changes in UE. In this way, operators can innovate inside their domain without having to depend on either over-the-top (OTT) service providers or UE vendors to support their innovations. More important, SDMN aims at making the forwarding substrate fully software-driven, enabling the creation of any network type on demand (as seen later in this article), and opens the mobile network to innovation through new service enablers.

Figure 2 illustrates the central elements of SDMN, contrasting it visually with today's EPC (Fig. 1). The key enablers in our architecture consist of the MobileFlow forwarding engine (MFFE) and MobileFlow controller (MFC). MFFEs are interconnected by an underlying IP/Ethernet transport network. Recall from Fig. 1 that in EPC, the SGW and PGW are user plane elements, while the mobility management entity (MME) is a control plane element. Figure 2 illustrates a new split that decouples mobile network control from all user plane elements. This way the (new) user plane (i.e., the MFFE) becomes simple, stable, and high-performing, while the control plane (i.e., the MFC and mobile network applications) can be implemented in a logically centralized manner. Forwarding in the MFFE can be fully defined in software, while the control software can flexibly steer user traffic to different service enablers (e.g., deep packet inspection [DPI], video caching, and optimization) that can be distributed throughout the mobile network, as shown in Fig. 2. This design, which follows in the tradition first explored in [6], permits us to facili-

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0235

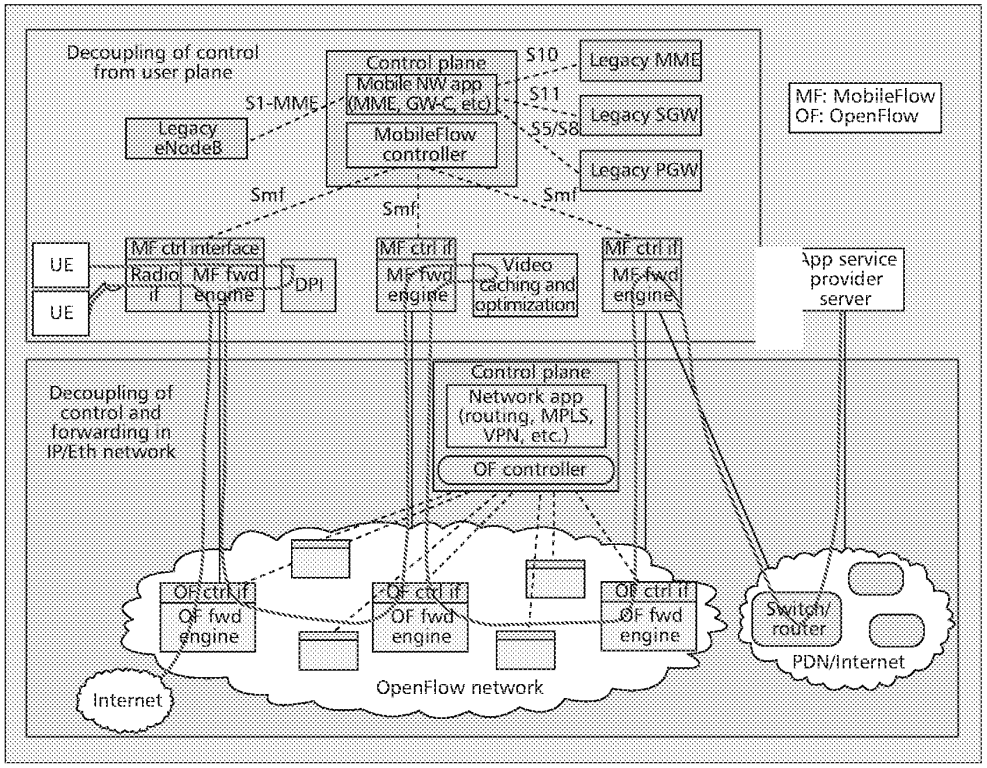


Figure 2. *Software defined mobile network.*

MFFEs include standard mobile network tunnel processing capabilities, such as, for instance, GTP-U and GRE encapsulation/ decapsulation facilities. Consequently, MFFEs, with the support of a MobileFlow controller, can be integrated with legacy EPC equipment.

tate swift network innovation through network function virtualization. For example, in this case, the carrier may opt to either reuse existing DPI equipment or employ virtualized functionality in a data center.

Note also that the figure illustrates two UE devices. The traffic from one of them is handled by a set of MFFEs, while the traffic from the second one is offloaded to the Internet primarily using the underlying OpenFlow transport network. Moreover, the depiction of virtualized network functions, such as DPI, and their interconnection with specific MMFEs is only indicative.

n/w app

Effectively, we introduce a new stratum based on a logically centralized MobileFlow controller that can steer forwarding. MFFEs employ network virtualization, are suitable for multitenant mobile networks, and can support advanced and security-related network functions. MFFEs are more complex than an OpenFlow switch, but much simpler than a router or a PGW, as the majority of control plane functionality is factored out to the MFC and the associated MobileFlow applications. In other words, significant parts of the functionality of today's EPC can be centralized into a carrier-grade control plane (upper part of Fig. 2). Figure 2 also depicts the OpenFlow-based decoupling of the IP/Ethernet transport network, contrasting it with the aforementioned mobile network control- and forwarding-plane decoupling. As can be surmised from Fig. 2, OpenFlow alone is not sufficient for implementing a carrier-grade mobile network user plane. We distinguish an SDMN from an OpenFlow-based network, as MFFEs are not switch-level equipment. MFFEs must support carrier-grade functionality such as network layer (L3) tunneling and flexible charging.

cloud device

network device

Our MobileFlow architecture can easily interact with legacy EPC network elements. In the control plane, the mobile network applications running on top of an MFC can function as an MME or the gateway control part of a physical box (indicated as GW-C), and thus can interoperate with other legacy MMEs, SGWs, or PGWs using the well established 3GPP-defined interfaces, as shown in Fig. 2. In the user plane, the MFFE can be interconnected with legacy EPC elements with a common IP/Ethernet transport network. The control messages can be transmitted via the MFFE in an in-band manner or via a dedicated out-of-band network.

Next, we look into the details of each of the architectural components of Fig. 2 and provide implementation insights based on our experience with an SDMN in our research laboratory.

network device

cloud app

MobileFlow Forwarding Engine

MFFEs include standard mobile network tunnel processing capabilities (e.g., GTP-U and GRE encapsulation/decapsulation facilities) [11]. Consequently, MFFEs, with the support of an MFC, can be integrated with legacy EPC equipment. An MFFE may also include the key functionality of a wireless access node, such as a radio interface to manage radio bearers. As illustrated in Fig. 2, an SDMN does not mandate any changes in UE. From a deployment perspective, one could envision wireless access MFFEs in parallel operation with existing eNodeBs, for example, as

eNodeB

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0236

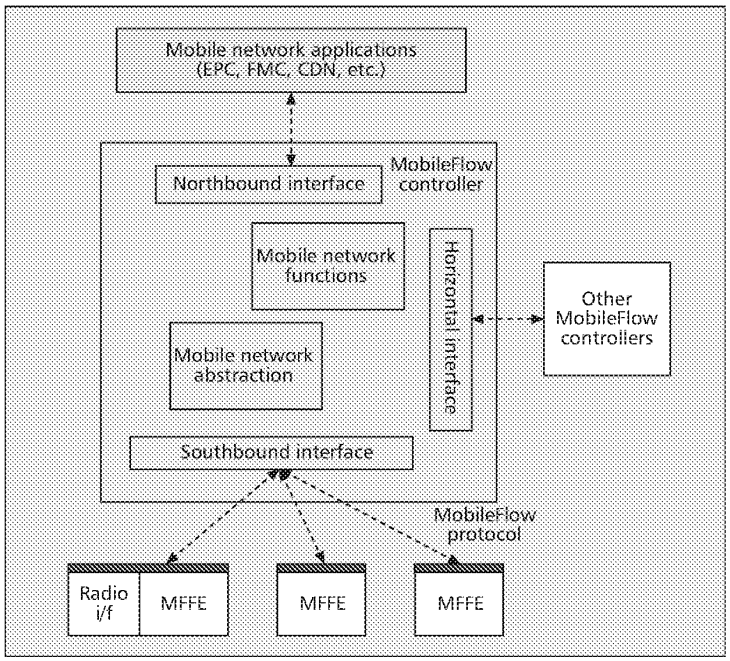


Figure 3. *MobileFlow controller.*

well as the incremental use of MFFEs in the core network as SGW/PGWs are retired.

Each MFFE communicates with an MFC through a lightweight protocol that implements the MobileFlow control interface (Smf). Due to space limitations we cannot delve into the details of this protocol, but suffice it to say that the MFFEs use the flow rules received from the MFC to encapsulate/decapsulate user traffic and/or forward packets to the transport network. Our current protocol implementation borrows from the flow-entry model advocated by the Open Networking Forum (ONF). Although we are still at a proof-of-concept stage, and more work is foreseen in the coming months, the performance we are observing in the laboratory is very promising, since the burdensome gateway control part is factored out of the user plane. Eventually, each MFFE featuring the lightweight Smf-driven control layer will be a high-performance, highly reliable network element that can handle a much larger number of flow entries than commodity OpenFlow switches do. This is because a future carrier mobile network should be able to handle very fine-grained policies, at the very minimum on par with what is now available for commercial PGWs. The comparative advantage is that with MFFEs, one can innovate much faster and, if so desired, can aim for different optimalities than what is possible today.

From an implementation perspective, MFFE can be based on a stripped-down mobile gateway that can handle network-layer (layer 3, L3) tunnel processing in the data plane with all other tunnel control processes moved up to the MFC. Alternatively, one can extend an OpenFlow switch to support the MobileFlow functions mentioned above. A hybrid model can be used as a combined forwarding element in the same IP/switch network level to facilitate 3GPP-IP convergence. For example, in Fig. 2, the MFFE can be combined with an OpenFlow switch, while the MFC can be combined with the OpenFlow controller.

MobileFlow Controller

Figure 3 zooms in on the MFC and illustrates its main functional blocks: the mobile network function, the mobile network abstraction, and the functional blocks corresponding to three interfaces. The MFC southbound interface controls MFFE operation. The horizontal interface is used for communication with other MFCs, and can be employed for intra- and trusted inter-domain cooperation. Operators with very large infrastructure typically segment their network according to administrative, regulatory, technical, or other reasons. In contrast to previous work [8, 9], our architecture specifies this interface from the beginning as we aim at very large carrier-grade deployments with high demands on reliability and availability that can scale out. Finally, the northbound interface facilitates fast network service and application development, as we discuss below.

The MFC relies on network-level abstraction. The corresponding functional block in Fig. 3 includes topology auto-discovery, topological resource view, network resource monitoring, and network resource virtualization. The network functions block includes tunnel processing, mobility anchoring, routing, and charging, among others. For instance, support for creating and establishing GTP-U tunnels is one of the network functions implemented in our testbed prototype, described later in this article. These abstractions correspond to high-level descriptions and are agnostic to the particular MFFE implementation. Thus, an operator can freely use MFFEs from different vendors. In addition, the operator can use MFFEs to adopt and deploy novel mobile network architectures that do not rely on tunneling.

A cornerstone of the SDN paradigm is the possibility of extending current and creating new functionality through programmatic APIs. An SDMN is compatible with advances in the OpenFlow-based ecosystem. Moreover, Fig. 3 illustrates the open interfaces and associated APIs introduced by the SDMN. Note that the SDMN controller does not have an explicit interface to legacy infrastructure. Any legacy interfacing can be handled efficiently by MobileFlow applications. The SDMN northbound interface, for example, can be used by MobileFlow applications to implement different network functionality. The MobileFlow stratum can be engaged in managing user traffic selectively depending on several criteria.

Returning to Fig. 2, recall that some traffic can be handled directly by the underlying OpenFlow-based transport network, while other traffic may require special treatment or support from advanced services. For example, mobility management, DPI, video optimization, and fine-grained QoS support can be selectively activated and handled at the MobileFlow stratum for premium traffic, while offloaded traffic can be handled by the transport network. We detail in the "Implementation" section later in this article

network and cloud devices

secure comm

network applns

communication

MFFE is a n/w device

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0237

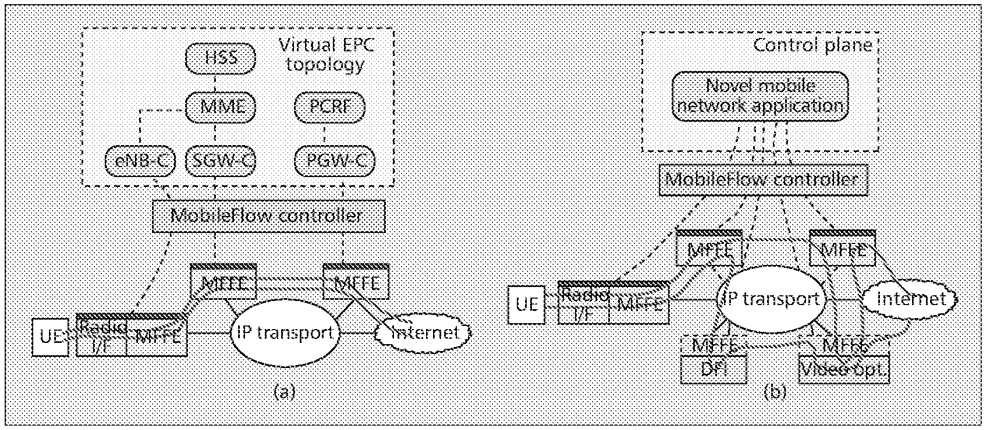


Figure 4. *MobileFlow application models.*

Supporting mobility in SDMN, whether based on existing standards by 3GPP and IETF or newly proposed by the research community, can be realized by introducing the corresponding applications above the northbound interface of the MFC.

how this is employed to realize different cellular architectures.

Mobile Network Applications

Mobile network applications can be developed by taking advantage of the SDMN northbound interface. Figure 4 illustrates the two models we explored while developing a fully EPS-compliant network based solely on an SDMN. The first model, illustrated in Fig. 4a, adopts a network function virtualization approach: a MobileFlow application implements the control part functionality of each of the EPS elements (e.g., eNodeB-C, SGW-C, PGW-C, S/PGW-C, and GGSN-C). Then each of these applications, which can run in a cloud computing environment similar to what is proposed in [10], interacts with the MFC to manage a respective MFFE, effectively implementing a virtual gateway. This model features a 1:1 mapping between the virtual gateway and MFFE via the MobileFlow controller. In other words, an EPC virtual topology is constructed in the MobileFlow stratum, as shown in Fig. 4a, according to the relationship between the corresponding reference architecture (Fig. 1) network elements (eNB-C, SGW-C, PGW-C, MME, PCRF, and HSS). Value-added functionality (e.g. video optimization) can be introduced by implementing the corresponding application and mapping it to an MFFE. This model is akin to the one introduced in [3], and although control may be centralized, application interaction is logically distributed.

The second model, shown in Fig. 4b, adopts a 1:*m* mapping between one all-encompassing MobileFlow application and several MFFEs. The application uses the SDMN northbound interface to access the substrate resource topology from the MobileFlow controller. In addition, it dictates the traffic paths and operations within the MFFE topology based on the network abstraction obtained from the MFC (Fig. 3). This model does not have to follow currently standardized interfaces and can therefore provide fertile ground for developing novel mobile network systems. For example, we can develop a virtual EPC gateway which aggregates functionality by factoring in that is currently distributed between several network elements, thus optimizing control signaling, while incorporating various SDMN-based service enablers, as shown in Fig. 4b.

Both models support *m*:1 mapping between MobileFlow applications and MFFEs, as each application belongs to a different control plane, therefore enabling multitenancy in a straightforward manner. Note that the first model is particularly suitable for integrating legacy infrastructure transparently into an SDMN system. The operator is free to employ the former model for one part of the network and the latter one for another part in which it can experiment with and readily deploy new network services. In fact, through our support for multitenancy we expect that it is possible to have both models in the same part of the network, although this is left as future work with respect to prototype implementation. Regardless of which model is preferred, SDMN enables flexible on-demand mapping between MobileFlow applications and their MFFE counterparts, as we explain in the following section.

Mobility Management

Supporting mobility in an SDMN, whether based on existing standards by 3GPP and IETF or newly proposed by the research community, can be realized by introducing the corresponding applications above the northbound interface of the MFC. For instance, as shown in Fig. 4a, where we take EPS as an illustrative case, we can use SDMN to implement the current generation of control functionality in a 3GPP mobile network. The wireless access MFFE with its eNB-C acts as an eNodeB. The MFFE with its associated SGW-C acts as the SGW, and the MFFE with its associated PGW-C acts as the PGW. In this case, mobility is managed by the virtualized but standard-adhering MME. The interfaces between MME and eNB-C/SGW-C are as defined by 3GPP (i.e., S1-MME and S11). The interface between SGW-C and PGW-C also follows the 3GPP-defined GTP-C part of S5/S8. The corresponding MFFEs handling the user traffic are controlled by their "-C" control counterparts with respect to processing GTP-U packets. The main difference from current deployments is that the -C virtualized function

distributed applications with network and cloud counterparts

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0238

We validated the SDMN architecture through research prototyping using COTS x86 based general-purpose servers and OpenFlow-based components. To showcase the flexibility of our architecture we developed the ODMN prototype, which demonstrates the core benefits of our MobileFlow approach.

translates the GTP signaling/context into flow rules and sends the rules via the MFC to each of the MFFEs involved in handling the particular flow. In other words, via the Smf interface we can control packet forwarding behavior, such as GTP tunnel encapsulation/decapsulation, quality of service (QoS) metering, forwarding, and so on, without having this type of control plane functionality in the boxes deployed in the field. The most challenging aspect here is how to split the functionality of an eNB, which is more complex than splitting the control and data planes in an S/P-GW. The interface between eNB-U and eNB-C consists of the forwarding control and wireless bearer control. The wireless bearer control is relevant to the radio handoffs required in mobility management processes, and at this stage of development this design aspect is part of our ongoing work. With respect to the prototype implementation we describe in the following section, the reader should keep in mind that the current code addresses only the split of S/PGW.

Looking forward, one can experiment on a small scale and then widely deploy novel mobility management solutions. In an SDMN this is done by upgrading or changing altogether the respective MobileFlow applications without replacing the deployed MFFEs. For example, Fig. 4b indicates that we can deploy a flat mobility management scheme which does not depend on S/PGW-like anchor nodes. This opens up opportunities for introducing simpler handover and gateway change processes, which today require lengthy standardization efforts. Earlier work on OpenPipes [6] and OpenRoads [8] explored a range of possibilities for optimized traffic steering through different service enablers such as video optimization, we discussed earlier in this article and illustrated explicitly by the different paths for the red and blue flows in Fig. 4b. Adopting a design based on locator/identifier split, including LISP, HIP, and others discussed in [15], can easily be accommodated. An SDMN aims to provide the same level of flexibility in handling mobility of users, flows, and content, but with carrier-grade performance. We are currently actively following the developments in the IETF Distributed Mobility Management (DMM) working group, and we plan to showcase in the future how SDMN can implement the novel solutions standardized there.

Implementation

We validated the SDMN architecture through research prototyping using COTS x86 based general-purpose servers and OpenFlow-based components. To showcase the flexibility of our architecture we developed the On-Demand Mobile Network (ODMN) prototype, which demonstrates the core benefits of our MobileFlow approach. In particular, we show that the MobileFlow stratum enables the software-based definition of different mobile networks (3G, 4G, FMC, and others) with different specifications (radio coverage, subscription numbers, PDP/bearer numbers, maximum bandwidth, etc.) according to administrator demand.

The ODMN is effectively a virtual network management system in which network resources include control and forwarding plane elements. We implement control plane network functions as mobile network applications (e.g., SGW-C, PGW-C, MME, and PCRF) and run them on general-purpose servers. Forwarding plane resources consist of radio access nodes and MFFEs. MFFEs may be implemented solely in software or using specific hardware platforms. By supporting virtualization for both control- and forwarding-plane resources, it is easy to create multiple virtual mobile networks with different types and specifications over the same hardware environment, as illustrated in Fig. 5. The same MFFE can be reused by different types of virtual networks. Each virtual mobile network can evolve and be upgraded solely by replacing the virtual machine running the corresponding MobileFlow application on the control plane servers — MFFEs remain unchanged.

flexible

Our testbed environment includes five servers in the control plane, as illustrated in the bottom left part of Fig. 6. Servers marked C1–C4 run multiple KVM virtual machines with different MFCs and applications (MME, SGW-C, PGW-C, GGSN-C, SGSN, etc.). A fifth server acts as the ODMN management server handling the creation, modification, and termination of virtual mobile networks. For example, Fig. 6 illustrates that we created three different networks. In the forwarding plane there are three radio access nodes (eNB1–2, NodeB) and three servers (F1–F3). eNB1 is a real Huawei eNodeB, while NodeB, radio network controller (RNC), and eNB2 are emulated using the Huawei Network Traffic Emulator (NTE). We developed each MFFE by extending the Open vSwitch code (available from www.openvswitch.org) to support GTP tunneling. All control plane and forwarding plane servers are interconnected by commercial off-the-shelf (COTS) LAN switches.

We also developed the ODMN testbed operator graphical user interface (Fig. 6, top left), which we refer to as the ODMN dashboard. This dashboard provides an overview of the physical resource topology and, more important, is used for all virtual resource allocations described next. The dashboard GUI also displays, over each network element icon, the virtual resource instances. Finally, the right side of Fig. 6 portrays different snapshots from the virtual networks we instantiate for the demonstration presented in this article.

Without loss of generality, we use ODMN to deploy three different mobile networks concurrently on the research testbed. The first network (marked in yellow in Fig. 6) is a standard 3G network composed of an NTE emulated RNC, a virtualized serving GPRS support node (SGSN), and a virtualized gateway GPRS support node (GGSN). As the color code indicates in the dashboard (Fig. 6, top left), the SGSN corresponds to an SGSN virtual machine (VM) on C1; the GGSN is composed of an MFFE VM on F3 and a corresponding GGSN-C VM on C4. The second virtual network (shown in red in Fig. 6) is a 4G network including the real eNodeB (eNB1), a virtual MME (C3), and a virtualized S/PGW. The virtual S/PGW comprises an MFFE VM on F1 and the corresponding S/PGW-C VM on C3. The third virtual network (in green) is

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0239

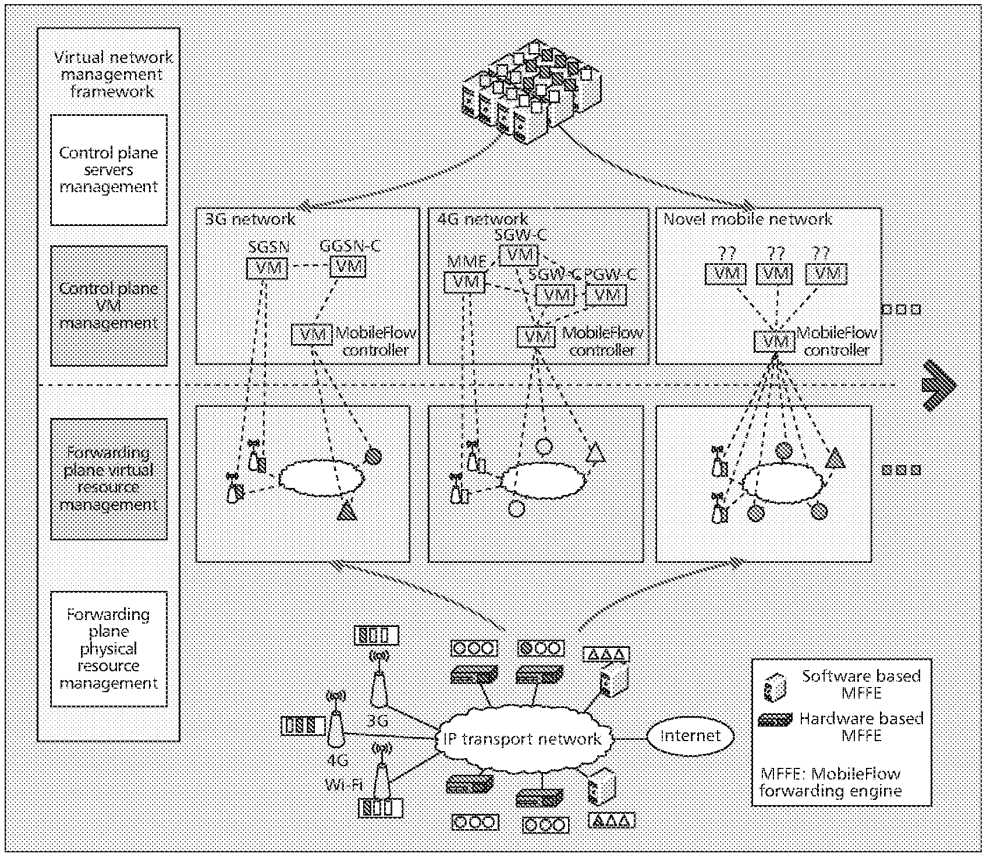


Figure 5. *On-demand mobile network.*

This essential functional test verified that the basic 3GPP attachment and bearer establishment procedures are correctly implemented using SDMN. However, the prototype does not fully support charging and policing; enhancing this aspect is part of our ongoing implementation work.

another 4G network including an NTE emulated eNodeB (eNB2), a virtual MME (C1), a virtual SGW, and a virtual PGW. The virtual SGW comprises an MFFE VM (F2) and the corresponding SGW-C VM on C2. The virtual PGW is composed of an MFFE VM on F3 and the corresponding PGW-C VM on C2.

In ODMN, a COTS LTE Android smart phone can successfully attach to the virtual 4G mobile network ("red network") via a real-world eNodeB (eNB1) through the standard 3GPP attachment procedure. Since the virtual S/PGW (corresponding to MFFE VM on F1) connects through the laboratory LAN to the Internet, the user of the LTE smartphone can surf the web through the ODMN-based virtual mobile network. This essential functional test verified that the basic 3GPP attachment and bearer establishment procedures are correctly implemented using an SDMN. However, the prototype does not fully support charging and policing; enhancing this aspect is part of our ongoing implementation work.

For the virtual 3G network ("yellow network" in Fig. 6) and virtual 4G network ("green network"), we verified their on-demand performance capability for the maximum bandwidth and number of PDP/bearer contexts. We specify the maximum bandwidth to 100 and 200 Mb/s for the virtual 3G and 4G networks, respectively. We specify the maximum number of PDP/bearer contexts to 200 for both virtual networks. We employed NTE to emulate 400 PDP/bearers attached to the two networks with 400 Mb/s maximum bandwidth, respectively, and verified that the ODMN can successfully match the maximum bandwidth and number of PDP/bearers as per each test specification. As illustrated above, virtualization is explicitly supported, while multitenancy is taken to a new level as mobile networks with different types and specifications can be deployed on demand.

Operator Benefits

By explicitly separating the transport stratum from the mobile network one, we introduce a new degree of freedom in our system, which is not present in an OpenFlow-based network. First, the two strata can be developed and deployed independently. This means that a mobile operator can continue to benefit from the reliability of current EPC equipment while rolling out a commodity OpenFlow-based transport network to reduce costs. Simultaneously, the operator can start to deploy MFFEs, which can interoperate with legacy equipment (e.g., via GTP/PMIP tunnels), and experiment with new services built on the software-defined part of the mobile network. Without the MobileFlow stratum, an operator would effectively have to discard all current equipment, including specialized middleboxes, and reintroduce the same functionality over the transport network applications space. This latter approach has many risks due to the

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0240

motivation

lack of carrier-grade mature OpenFlow products. Instead, our architecture enables operators to choose the best of both worlds and migrate to a fully software-defined network on their own initiative and pace. Furthermore, the architecture is forward-looking, so the operator can also replace the transport stratum without losing its investment in the MobileFlow stratum. This is of particular interest, for example, when considering fixed mobile convergence (FMC).

As our research prototype proof-of-concept results indicate, SDMN enables operators to use infrastructure resources to orchestrate on demand the creation of various mobile network pipes based on different mobile architectures (3G, 4G, FMC, etc.) using the appropriate MobileFlow applications. While doing so, the operator can manage traffic steering for value-added services, enabling new possibilities for improving the end-user quality of experience as well as optimizing the cost of handling lesser-value traffic flows. An example of the latter includes offloading selected traffic to the closest egress at the transport network level without having to traverse the mobile core, as illustrated in Fig. 2. Since the MFC maintains a complete view of the infrastructure network topology, MobileFlow applications can be used to orchastrate various service chains, thus delivering a personalized user experience.

Last but certainly not least, SDMN enables speedy and smooth mobile network evolution and differentiation. As the example in Fig. 5 illustrates, an operator can start using SDMN with a 3G network in place and evolve it into a 4G network by allocating more resources as the rollout proceeds. Changing important mechanisms at the core is also decoupled from evolution at the edge. Today this would require a technology swap, which is costly, time-consuming and may introduce competitive deployment races that can hurt operator profits. Through a phased approach that can follow each operator's own pace, future SDMN-based networks can be deployed and operated. Moreover, within the same domain, a carrier can allocate a dynamic

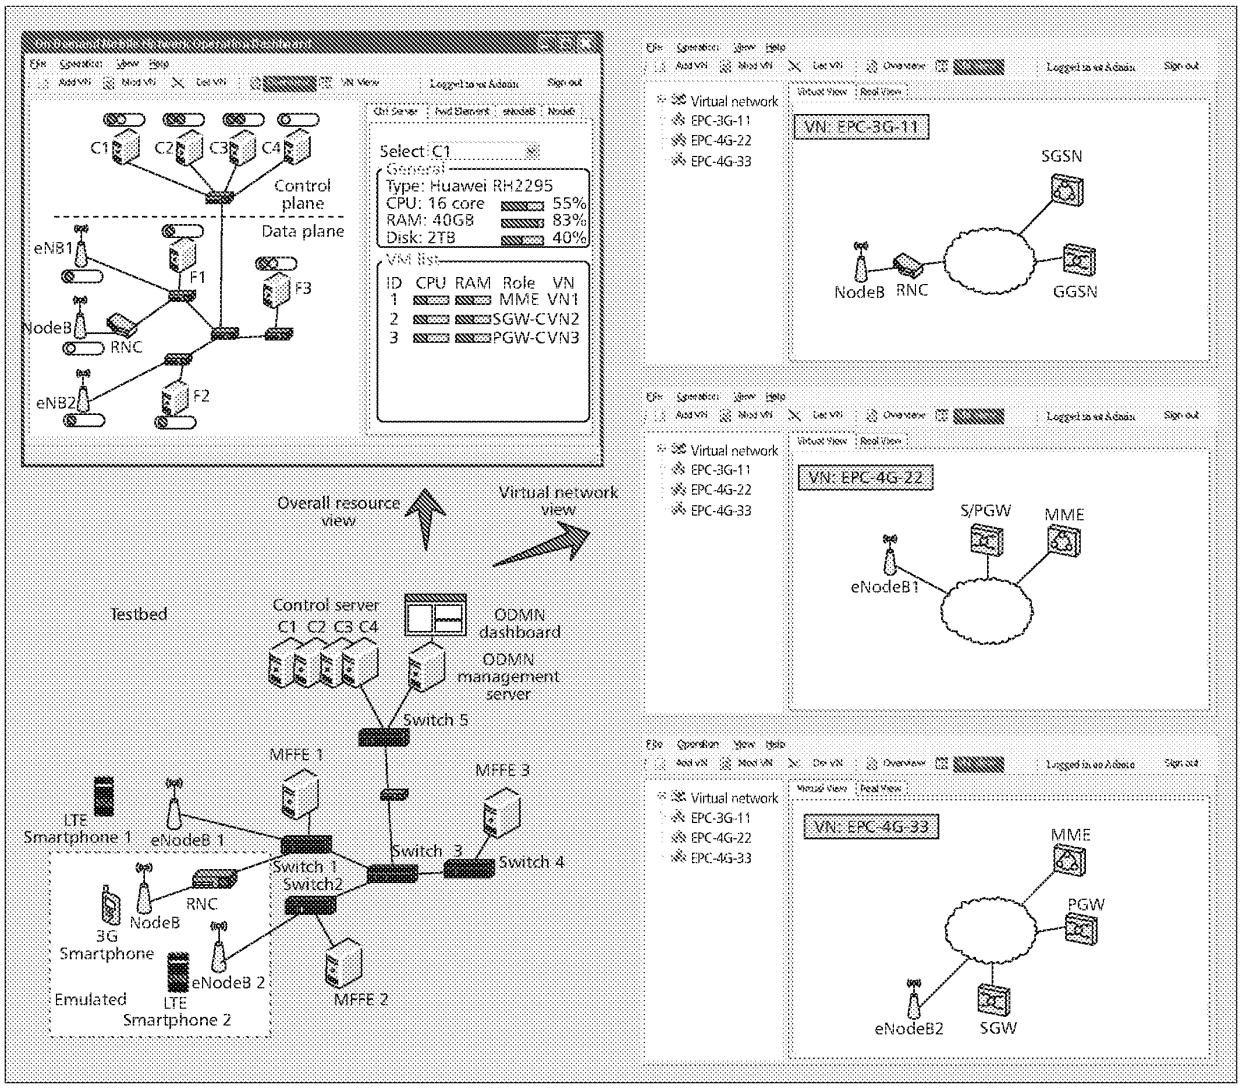


Figure 6. *ODMN prototype testbed and dashboard user interface.*

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0241

mix of control and user plane resources from the virtualized network infrastructure to address different pain points as they emerge (e.g., signaling storms or spikes in content distribution during popular events) on much shorter timescales than today. Within this context, operators may procure (or develop) different sets of MobileFlow applications based on the subscriber groups they address and thus further differentiate themselves from their competition.

Conclusion

We introduce SDMN, an architecture for future carrier networks that builds on the decoupling of data and control in the mobile network user plane and a new MobileFlow stratum, which can significantly increase the operator innovation potential. The SDMN open interfaces and APIs foster service innovation, increasing the capability of the operator to roll out new network features while reducing time to market for new services. Our proof-of-concept prototype demonstrates that by employing the SDMN architecture a carrier can configure on-demand and in detail the network architecture, radio coverage, gateway location, PDP/bearers number, maximum throughput, and so on. Moreover, as we have seen, an SDMN operator can orchestrate new service chains and evolve its network by updating or replacing MobileFlow applications.

Our ODMN testbed implementation validates the programmability and flexibility of the SDMN approach. That said, many challenges lie ahead as we focus on improving performance, ease of resource configuration and management, scalability, and code maturity. ODMN provides the foundation for further research and experimentation in this area. We expect that SDMNs can catalyze operator innovations in a range of areas, from speedy introduction of new services to improved monitoring and management of network resources and services, to controlling CAPEX/OPEX, to personalized handling of subscriber traffic to yield more revenue and reduce churn. With SDMN, operators can differentiate their offerings at an unprecedented scale, while protecting their current investment in traditional cellular equipment.

Acknowledgment

We are grateful to the anonymous reviewers for their constructive comments and suggestions, which helped us to hone the core message of this article. We thank Dr. Jiang Li and Hua Huang for their comments and suggestions, which helped improve the overall quality of this work. The views expressed herein are solely those of the authors and do not necessarily represent the views of Huawei Technologies.

References

[1] A. Greenhalgh *et al.*, "Flow Processing and the Rise of Commodity Network Hardware," *SIGCOMM CCR*, vol. 39, no. 2, ACM, 2009, pp. 20–26.
[2] N. McKeown *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *SIGCOMM CCR*, vol. 38, no. 2, ACM, 2008, pp. 69–74.
[3] M. R. Nascimento *et al.*, "Virtual Routers as A Service: the RouteFlow Approach Leveraging Software-Defined Networks," *Proc. CFI*, Seoul, Korea. ACM, 2011.
[4] C. J. Sher Decusatis *et al.*, "Communication within Clouds: Open Standards and Proprietary Protocols for Data Center Networking," *IEEE Commun. Mag.*, vol. 50, no. 9, Sept. 2012.
[5] M. Bari *et al.*, "Data Center Network Virtualization: A Survey," *IEEE Commun. Surveys & Tutorials*, vol. PP, no. 99 (early access article).
[6] G. Gibb *et al.*, "OpenPipes: Prototyping high-speed networking systems," *Proc. SIGCOMM (Demo)*, Barcelona, Spain, 2009.
[7] T. Koponen *et al.*, "Onix: A Distributed Control Platform for Large-Scale Production Networks," *Proc. OSDI*, Vancouver, BC, Canada, 2010.
[8] K.-K. Yap *et al.*, "Blueprint for Introducing Innovation into Wireless Mobile Networks," *Proc. SIGCOMM VISA Wksp.*, New Delhi, India, 2010.
[9] E. Li *et al.*., "Toward Software-Defined Cellular Networks," *Proc. EWSDN*, Darmstadt, Germany, 2012.
[10] J. Kempf *et al.*, "Moving the Mobile Evolved Packet Core to the Cloud," *Proc. WiMob*, Barcelona, Spain, 2012.
[11] I. Ali *et al.*, "Network-based Mobility Management in the Evolved 3GPP Core Network," *IEEE Commun. Mag.*, vol. 47, no. 2, Feb. 2009.
[12] H. Ekström, "QoS Control in the 3GPP Evolved Packet System," *IEEE Commun. Mag.*, vol. 47, no. 2, Feb. 2009.
[13] J.-J. Pastor Balbás *et al.*, "Policy and Charging Control in the Evolved Packet System," *IEEE Commun. Mag.*, vol. 47, no. 2, Feb. 2009.
[14] K. Pentikousis *et al.*, "Design Considerations for Mobility Management in Future Infrastructure Networks," *Proc. ITU TELECOM WORLD Technical Symp.*, Geneva, Switzerland, 2011.
[15] B. Sousa *et al.*, "Multihoming Management for Future Networks," *ACM/Springer Mobile Networks and Applications*, vol. 16, no. 4, Aug. 2011, pp. 505–17.

Biographies

Kostas Pentikousis (k.pentikousis@huawei.com) is a senior research engineer at Huawei Technologies, Berlin, Germany. He earned his Bachelor's degree in informatics from Aristotle University of Thessaloniki, and his Master's and doctoral degrees in computer science from Stony Brook University. At Huawei he worked on 3GPP EPC research topics beyond Rel. 12, carrier-grade SDN, and NFV.

Yan Wang (jason.wangyan@huawei.com) is a senior research engineer at Huawei Technologies, Shanghai, China. He earned his doctoral degree in communication and information systems from the State Key Laboratory on Advanced Optical Communication Systems and Networks, Shanghai Jiao Tong University (SJTU), China. At Huawei he worked as the leader of the Software Defined Mobile Network (SDMN) group. His research interests include 3GPP packet core research topics, carrier-grade SDN, and NFV.

Weihua Hu (huweihua@huawei.com) is a senior research engineer at Huawei Technologies in Shanghai, China. He earned his Bachelor's and Master's degrees in control theory and control project from SJTU. At Huawei he worked on 3GPP packet core research topics, carrier-grade SDN, and NFV.

Our ODMN testbed implementation validates the programmability and flexibility of the SDMN approach. That said, many challenges lie ahead as we focus on improving performance, ease of resource configuration and management, scalability, and code maturity.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0242

PTO/AIA/26 (04-14)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT	Docket Number (Optional) TAAG001CONT1

In re Application of: Taaghol

Application No.: 15/836,824

Filed: 12/9/2017

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

The applicant, Pouya Taaghol, owner of 100 percent interest in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of **prior patent** No. 9843624 as the term of said **prior patent** is presently shortened by any terminal disclaimer. The applicant hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the **prior patent** are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the applicant does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the **prior patent**, "as the term of said **prior patent** is presently shortened by any terminal disclaimer," in the event that said **prior patent** later:

expires for failure to pay a maintenance fee;
is held unenforceable;
is found invalid by a court of competent jurisdiction;
is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
has all claims canceled by a reexamination certificate;
is reissued; or
is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

Check either box 1 or 2 below, if appropriate.

1. ☐ The undersigned is the applicant. If the applicant is an assignee, the undersigned is authorized to act on behalf of the assignee.

I hereby acknowledge that any willful false statements made are punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

2. ☑ The undersigned is an attorney or agent of record. Reg. No. 40,394

/Christopher J. Horgan/ — Signature
October 16, 2019 — Date

Christopher J. Horgan — Typed or printed name

Attorney for Applicant — Title
408-687-8306 — Telephone Number

☑ Terminal disclaimer fee under 37 CFR 1.20(d) included.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

This collection of information is required by 37 CFR 1.321. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Electronic Patent Application Fee Transmittal

Application Number:	15836824
Filing Date:	09-Dec-2017
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Filer:	Christopher James Horgan
Attorney Docket Number:	TAAG001CONT1

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension - 1 month with $0 paid	2251	1	100	100
Miscellaneous:				
STATUTORY OR TERMINAL DISCLAIMER	2814	1	160	160
		Total in USD ($)		**260**

Electronic Acknowledgement Receipt

EFS ID:	37568833
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	25-OCT-2019
Filing Date:	09-DEC-2017
Time Stamp:	17:18:04
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$260
RAM confirmation Number	E20190OH24521608
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Amendment/Req. Reconsideration-After Non-Final Reject	TAAG001CONT1_RESPONSE_.pdf	104191 c5c33448506c4c6eae5fb2234920b95f21dd34b2	no	14
Warnings:					
Information:					
2	Terminal Disclaimer Filed	TAAG001CONT1_TD.pdf	142282 cf27d5666d7d6622e46e98a5520e391faa1a487d	no	1
Warnings:					
Information:					
3	Fee Worksheet (SB06)	fee-info.pdf	32741 01abc58debbf39b4960c76bde8fe9b74e2d0033b	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	279214		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.
National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.
New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Taaghol
Serial No.: 15/836,824
12/09/2017
Attorney Docket No.: TAAG001CONT1

Examiner: Krishnan, Rajeshwari
Art Unit: 2441 Filed:

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

Mail Stop Amendment Commissioner for Patents PO Box 1450
Alexandria, VA 22313-1450

Dear Sir:

AMENDMENT

Applicants offer the following election in response to the Office Action (hereinafter "Office Action") mailed on June 25, 2019.

Amendments to the Specification being on page 2 of this paper.
Amendments to the Claims begin on page 5 of this paper.
Remarks/Arguments begin on page 11 of this paper.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0248

In the specification:

Please make the following correction In paragraph [0025] in the Table.

[0025] …

fxManager	FleXible Manager that ~~manager~~ manages the fxOS lifecycle (provisioning, usage, de-provisioning)

Please replace paragraph [0068] with the following amended paragraph.

[0068] The flexible Software Development Kit (fxSDK) may include a development environment and other tools to facilitate development, testing, debugging and verifications of applications for dSDN environments. The fxSDK is unique in several aspects exemplified as follows. First, since the dSDN applications may have a device component (fxDevice 302) and cloud component (fxCloud 304), the fxSDK allows the developer to develop the components together to simplify the development and testing. Second, the dRMS offers various system level APIs with remote accessibility. In other words, the APIs in the fxDevice 302 could be accessed via the applications in the fxCloud (fxCloudApps) and vice versa. The fxSDK simplifies the usage of system created APIs and application created APIs. Third, it is essential for the dSDN to ensure quality and security of applications in the network. The fxSDK may be accompanied with an extensive Risk Analysis Utility (RAU) which verifies all the APIs that are being used by the application and based on its potential danger to the network, it would show the risk analysis and may make specific suggestions to reduce the risk. Once completed, the risk analysis result is included in the dAP for posting to the fxStore and consumed by the fxManager user (the admin). Fourth, the fxSDK would also ensure that the applications do not exceed their planned memory and resource usage. Fifth, the Integrated Development Environment (IDE) may provide

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0249

various tools to actively demonstrate errors, defined APIs, security dangers, and excessive resource usage. Sixth, the IDE could also provide a flexible Simulator (fxSimulator) for the developer test the application in a network setup in presence of other network elements.

Please replace paragraph [0084] with the following amended paragraph.

[0084] The Virtual Fabric (fxVF) provides an abstraction layer for ~~application~~ applications to communicate with each other whether they are in the fxDevice or fxCloud. Various frameworks and services may use the fxVF service. Figure 15 illustrates fxVF 1502a, 1503a, 1504a in the fxDevices 1502 and 1503 and fxCloud 1504. The fxVF 1502a, 1503a, 1504a may use Extensible Messaging and Presence Protocol (XMPP) for this messaging as the default protocol. It is important that the developers may use their own communication protocols between the fxDeviceApps 1502b, 1503b and the associated fxCloudApp 1504b. fxVF 1502a, 1503a, 1504a provide a secure routing mechanism.

Please replace paragraph [0088] with the following amended paragraph.

[0088] An app (fxDeviceApp or fxCloudApp) could expose APIs to be used by the other apps. The inter-app communication is enabled by the fxVF where the policy and security provisions are enforced. The APIs exposed may be RESTful (representational state transfer) and could travel across physical network elements. Since the fxManager defines fxVF policies, the administrator could ultimately specify which APIs between which apps could communicate with each other. Figure 16 demonstrates the logical interfaces between the platform resources and applications. As shown in Figure 16, the distributed Resources Service (dRS) 1602, 1603 comprises software agents residing in the fxDevice 1604 and the fxCloud 1606. These dRS agents manage the access to the platform ~~resolutions~~ resources (including potential configuration conflict resolution) as well as facilitating inter-app APIs. The dRS 1602, 1603 allows or disallows access to platform resource or inter-app communication according to the policies defined by the network administrator using the fxManager. In summary, the dRS 1602, 1603 provides services to application developers for:

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0250

Exposing APIs to other applications;

Configuring and managing platform

resource;

Policy enforcement and authorization of applications access to platform

resource and other app's APIs; and

Policy conflict resolution.

Please replace paragraph [0096] with the following amended paragraph.

[0096] It is important for the target network element to verify the authenticity and integrity of loaded software prior to install. The uniquely defined security keys of the application developers sign the applications. The following rules may be followed to ensure software securityverification:

Only the verified OS/Firmware software can install on verified hardware;

Only the verified OS/Firmware software can run on verified hardware (continuous verification of essential parts of the software);

Only the verified fxDeviceApp software can install on verified hardware;

Only the verified fxDeviceApp software can run on verified hardware (continuous verification of essential parts of the software).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0251

In the Claims:

1. (Currently Amended) A system for processing data packets in a network, the system comprising:

a programmable network device hosting a plurality of first network applications;

a programmable cloud device hosting a plurality of second network applications[;], wherein at least one of the plurality of first network applications in the programmable network device and at least one of the plurality of second network applications in the programmable cloud device are in secure communication with each other through a virtual fabric to form a distributed application;

wherein the programmable network device and programmable cloud device are each powered by a sandboxing operating system which facilitates deployment of the plurality of first and second network applications ~~and facilitates upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device~~;

an application management portal coupled to the programmable network device and programmable cloud device and capable of managing upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device;

wherein the ~~programmable network device~~ application management portal verifies authenticity of the upgrades to the plurality of first network applications and the ~~programmable cloud device verifies authenticity of the upgrades to the~~ plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

an application repository coupled to the application management portal and capable of storing distributed applications ~~which have been tested~~ for installation in the programmable network device and programmable cloud

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0252

device; and

~~an application management portal coupled to the programmable network device, programmable cloud device and the application repository,~~ wherein the application management portal <u>further</u> manages usage of the distributed applications on the programmable network device and programmable cloud device.

2. (Previously Presented) The system of claim 1, wherein each of the programmable network device and programmable cloud device include an access network interface unit to send and receive communications and a processor with a memory associated with the network interface unit.

3. (Canceled)

4. (Canceled)

5. (Previously Presented) The system of claim 1, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines, the network functions consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).

6. (Canceled)

7. (Previously Presented) The system of claim 1, further comprising:

a second programmable network device which includes the at least one of the plurality of first network applications which is in the first programmable network device.

8. (Canceled)

9. (Canceled)

10. (Canceled)

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0253

11. (Currently Amended) The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a distributed content service wherein ~~contents~~ content generated by the at least one of the plurality of first network applications may be shared though the virtual fabric with the at least one of a plurality of second network applications.

12. (Currently Amended) The system of claim 1, wherein the programmable network device and programmable cloud device further comprise:

a distributed notification service wherein notice of a particular event is sent through the virtual fabric from at least one of the plurality of first network applications to at least one of the plurality of second network applications.

13. (Currently Amended) The system of claim [10] 1, ~~wherein the DRS further includes a load controller; and wherein the~~ further comprising a load controller [is] adapted to monitor loads on the programmable network device and the programmable cloud device ~~at least one of the plurality of first network applications and at least one of the plurality of second network applications~~ and ~~effect~~ affect change in accordance with thresholds received from the application management portal.

14. (Previously Presented) The system of claim 1, wherein the programmable network device, at least one of the plurality of first network applications, the programmable cloud device, and at least one of the plurality of second network applications each have unique identities.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0254

15. (Currently Amended) A system for processing data packets, the system comprising:

a programmable network device processing data flows hosting at least one of a plurality of first network applications;

a programmable cloud device having a plurality of virtual machines wherein at least one of the virtual machines hosts at least one of a plurality of second network applications[;], wherein the at least one of the plurality of first and the at least one of the plurality of second network applications are in secure communication through a virtual fabric to form a distributed application;

wherein the programmable network device and programmable cloud device are each powered by a sandboxing operation system which facilitates deployment of the plurality of first and second network applications ~~and facilitates upgrade of the plurality of first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device~~;

an application management portal coupled to the programmable network device and programmable cloud device and capable of managing upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device;

wherein the ~~programmable network device~~ application management portal verifies authenticity of the upgrades to the plurality of first network applications and the ~~programmable cloud device verifies the authenticity of the upgrades to the~~ plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

an application repository storing distributed applications ~~which have been tested~~ for installation in the programmable network device and programmable cloud device; and

~~an application management portal coupled to the programmable~~

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0255

~~network device, programmable cloud device and the application repository,~~ wherein the application management portal further manages provisioning, usage and de-provisioning of the distributed applications on the programmable network device and programmable cloud device.

16. (Previously Presented) The system of claim 15, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines, the network functions consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).

17. (Currently Amended) A system for processing data packets in a network, the system comprising:

a programmable network device processing data flows, hosting at least one of a plurality of first network applications;

a programmable cloud device partitioned into a plurality of zones, wherein at least one of the plurality of zones hosts at least one of a plurality of second network applications[;], wherein the at least one of the plurality of first network applications and at least one of the plurality of second network applications are in secure communication through a virtual fabric to form a distributed application;

wherein the programmable network device and programmable cloud device are each powered by a sandboxing operation system which facilitates deployment of the plurality of first and second network applications ~~and facilitates upgrade of the plurality of first and second network applications with substantially no interruption to operation of the programmable network device and programmable cloud device~~;

an application management portal coupled to the programmable network device and programmable cloud device and capable of managing upgrades of the first and second network applications with substantially no interruption to operation of the programmable network device and

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0256

programmable cloud device

wherein the ~~programmable network device~~ application management portal verifies authenticity and integrity of upgrades to the plurality of first network applications and the ~~programmable cloud device verifies the authenticity of the upgrades to the~~ plurality of second network applications and wherein the verification is based on unique security keys associated with each of the plurality of first and second network applications;

an application repository storing distributed applications which have been tested for installation in the programmable network device and programmable cloud device; and

~~an application management portal coupled to the programmable network device, programmable cloud device and the application repository,~~ wherein the application management portal further manages usage of the distributed applications on the programmable network device and programmable cloud device.

18. (Previously Presented) The system of claim 17, wherein the programmable network device is a Long Term Evolution (LTE) evolved NodeB (eNB) and the programmable cloud device hosts LTE core network functions in separate virtual machines, the network functions consisting of: Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0257

REMARKS

Claims 1-2, 5, 7 and 11-18 are currently pending. By this Amendment, Applicant has canceled claim 3, 4, 6 and 8-10 and amended claims 1, 11, 12, 13, 15 and 17.

OBJECTIONS TO THE SPECIFICATION

The specification was objected to for several spelling errors. Applicant has made the appropriate corrections.

CLAIM OBJECTIONS

The Examiner has objected to claims 1, 13, 15 and 17. Applicant has made the appropriate corrections.

CLAIM INTERPRETATION – CONTINGENT CLAUSES NOT EVALUATED

The Examiner has indicated that limitations in several of the claims will be given no patentable weight. The Applicant has addressed these issues in the following ways. Firstly, the Applicant has added the "virtual fabric limitation to claims 1, 15 and 17. Secondly, the Applicant has clariid in the claims that the application management portal manages the upgrades. Third, Applicant has deleted the limitation regarding the testing of the distributed applications. Fourthly, claims 3 and 4 have been canceled. Fifthly, Applicant has added in claim 11 that the distributed content service shares the content through the virtual fabric. Sixthly, Applicant has added in claim 12 that the notice is sent through the virtual fabric.

CLAIM INTERPRETATION – 35 USC 112(f)

The Examiner has indicated that limitations in several of the claims will be interpreted as means plus function claims. The Applicant has addressed these issues in the following ways. Firstly, Applicant has canceled claims 6, 8 and 10. Secondly, Applicant acknowledges that claim 13 may be treated as a means plus function claim.

CLAIM REJECTIONS – 35 USC 112(a),112(f)

Claims 6, 8 and 10 are rejected under 35 USC 112(a) for failing to comply

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0258

with the written description requirement. Accordingly, claims 6, 8 and 10 have been deleted.

CLAIM REJECTIONS – 35 USC 112(b),112(f)

Claims 6, 8 and 10 are rejected under 35 USC 112. Accordingly, claims 6, 8 and 10 have been deleted.

CLAIM REJECTIONS – 35 USC 112(a)

With regard to independent claims 1, 15 and 17, on page 17 of the Office Action the Examiner states that "the disclosure fails to sufficiently identify how the fxOS performs the function of the hot upgrade with substantially no interruption". The Examiner's attention is called to paragraphs [0060] and [0077] to [0083] of the specification for a discussion of implementation of the hot upgrading.

Also, with regard to claims 1, 15 and 17 on page 17 of the Office Action, the Examiner asserts that the application management portal verifies the application. Applicant has amended claims 1, 15 and 17 accordingly.

Further, with regard to claims 1, 15 and 17 on page 18 of the Office Action, the Examiner objects to the use of the testing language in the claims. Accordingly, this language has been deleted.

With regard to claim 11, on page 18 of the Office Action, the Examiner objects to the language of the claim. Applicant has added language to the claims to make it clear that the information is shared through the virtual fabric.

With regard to claim 12, on page 19 of the Office Action, the Examiner objects to the language of the claim. Applicant has language to the claims to make it clear that the notice is sent through the virtual fabric.

CLAIM REJECTIONS – 35 USC 112(b)

On page 19 of the Office Action, the Examiner to the antecedent basis in claim 9. Applicant has canceled claim 9 by this amendment.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0259

DOUBLE PATENTING REJECTION

Claims 1-18 are rejected on the ground of nonstatutory double patenting as being unpatentable over Claims 1-14, 18 and 19 of U.S. Patent No. 9,843,624. Accordingly, Applicant has filed a Terminal Disclaimer with this response.

REJECTIONS UNDER 35 USC 103

Claims 1, 2, and 5-18 are rejection under 35 USC 103 as being unpatentable over Pentikousis, "MobileFlow: Toward Software-Defined Mobile Networks", IEEE Communications Magazine, July 2013 (hereinafter, "Pentikousis") in view of U.S. Patent Application Serial No. 20130223442 to Narayanan (hereinafter, "Narayanan") and U.S. Patent Application Serial No.20140317293 to Shatzkamer (hereinafter, "Shatzkamer"). Claim 4 is rejected under Pentikousis in view of Narayana and Shatzkamer and further in view of U.S. Patent Application No. 20140052976 to Marino (hereinafter "Marino"). Claim 5 is rejected under Pentikousis in view of Narayana and Shatzkamer and further in view of U.S. Patent Application No. 20130263104 to Baset (hereinafter "Baset").

Applicant notes that all of the rejections use the Pentikousis reference as the base reference. Applicant filed U.S. Patent Provisional Application No. 61/834,807 on June 13, 2013 upon which this application ultimately depends and which predates the Pentikousis reference published in July 2013. Therefore Applicant submits that the outstanding 35 USC 103 rejections are overcome by this fact.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0260

CONCLUSION

Therefore it is respectfully submitted that claims 1-2, 5, 7 and 11-18 are allowable over the prior art. No new matter has been added in the new claims and/or limitations presented in this Amendment.

Respectfully submitted,

ROARK IP

By: /Christopher J. Horgan/

Christopher J. Horgan
Registration No. 40,394
1438 Dahlia Loop
San Jose, CA 95126
Telephone: (408) 687-8306

Date: October 25, 2019

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0261

Application Number * 15/836,824 *	Application/Control No. 15/836,824	Applicant(s)/Patent under Reexamination Taaghol et al.	
	Examiner CHAN, WING F	Art Unit 2441	

Document Code - DISQ	Internal Document - DO NOT MAIL

TERMINAL DISCLAIMER	☐ APPROVED	☑ DISAPPROVED
Date Filed: **25 October 2019**	**This patent is subject to a Terminal Disclaimer**	

Approved/Disapproved by:

/CHARRISSA S BROWN/

Technology Center: PLRC

Telephone: (571)272-1558

The terminal disclaimer identifies a party who is not the applicant (only for applications filed on or after September 16, 2012; See FP 14.26.10):

For cases filed on/after 9/16/12, 37 CFR 1.321 specifies that only the applicant can disclaim, and the terminal disclaimer must specify the extent of the applicant's ownership.

Below is what needs to be done to correct the defects:

A request under 37 CFR 1.46(c) to change the applicant needs to be filed, which is (1) a request, signed by a 1.33 (b) party, (2) a corrected ADS (37 CFR 1.76(c)) that identifies the "new" applicant in the applicant information, and is underlined since it is new, and (3) a 3.73(c) statement showing chain of title to the new applicant. Along with the § 1.46(c) request we need a POA that gives power to the attorney who is signing the TD, along with another copy of the TD, or a TD that is signed by the applicant.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0262

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

112449 7590 01/21/2020

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

EXAMINER
HIGA, BRENDAN Y

ART UNIT	PAPER NUMBER
2441	

MAIL DATE	DELIVERY MODE
01/21/2020	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 15/836,824	Applicant(s) Taaghol et al.	
	Examiner BRENDAN HIGA	Art Unit 2441	AIA (FITF) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☑ Responsive to communication(s) filed on 10/25/2019.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☑ This action is **FINAL.** 2b) ☐ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5) ☑ Claim(s) 1-2,5,7 and 11-18 is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6) ☐ Claim(s) _____ is/are allowed.

7) ☑ Claim(s) 1-2,5,7 and 11-18 is/are rejected.

8) ☐ Claim(s) _____ is/are objected to.

9) ☐ Claim(s) _____ are subject to restriction and/or election requirement

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to **PPHfeedback@uspto.gov.**

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☑ Notice of References Cited (PTO-892)

2) ☐ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date _____.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date _____.

4) ☐ Other: _____.

DETAILED ACTION

This case has been reassigned to a new Examiner, see below for updated Examiner contact information.

This Office action is in response to Applicant's amendment and request for reconsideration filed on October 25, 2019.

Claims 1-2, 5, 7, and 11-18 are pending.

Response to Arguments

Applicant's arguments, see remarks filed October 25, 2019, with respect to the previous rejection of claims 1-2, 5, 7, and 11-18 under 35 USC 112(a) and/or 112(b), and 103 have been fully considered and are persuasive. The rejection of claims 1-2, 5, 7, and 11-18 under 35 USC 112(a), 112(b), and 103 have been withdrawn.

Terminal Disclaimer

The terminal disclaimer does not comply with 37 CFR 1.321 because:

This application was filed on or after September 16, 2012. For cases filed on or after September 16, 2012, 37 CFR 1.321 specifies that only the applicant can disclaim, and the terminal disclaimer must specify the extent of the applicant's ownership.

Below is what needs to be done to correct the defects:

A request under 37 CFR 1.46(c) to change the applicant needs to be filed, which is (1) a request, signed by a 1.33 (b) party, (2) a corrected ADS (37 CFR 1.76(c)) that identifies the "new" applicant in the applicant information, and is underlined since it is

new, and (3) a 3.73(c) statement showing chain of title to the new applicant. Along with the § 1.46(c) request we need a POA that gives power to the attorney who is signing the TD, along with another copy of the TD, or a TD that is signed by the applicant.

Double Patenting

The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory double patenting rejection is appropriate where the claims at issue are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); In re Vogel, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on a nonstatutory double patenting ground provided the reference application or patent either is shown to be commonly owned with this application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement. See MPEP § 717.02 for applications subject to examination under the first inventor to file provisions

of the AIA as explained in MPEP § 2159. See MPEP §§ 706.02(l)(1) - 706.02(l)(3) for applications not subject to examination under the first inventor to file provisions of the AIA. A terminal disclaimer must be signed in compliance with 37 CFR 1.321(b).

The USPTO Internet website contains terminal disclaimer forms which may be used. Please visit www.uspto.gov/forms/. The filing date of the application in which the form is filed determines what form (e.g., PTO/SB/25, PTO/SB/26, PTO/AIA/25, or PTO/AIA/26) should be used. A web-based eTerminal Disclaimer may be filled out completely online using web-screens. An eTerminal Disclaimer that meets all requirements is auto-processed and approved immediately upon submission. For more information about eTerminal Disclaimers, refer to http://www.uspto.gov/patents/process/file/efs/guidance/eTD-info-I.jsp.

Claims 1, 2, 5, 7, and 11-18 are rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1-21 of **US Patent No. 9,843,624** (624' Patent). The current application merely broadens the scope of the 624' Patent, as such the current Application is anticipated by the narrower claim scope of the 624' Patent.

Allowable Subject Matter

Claims 1, 2, 5, 7, and 11-18 are allowable over the prior art.

Conclusion

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure (see PTO 892).

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to BRENDAN HIGA whose telephone number is (571)272-5823. The examiner can normally be reached on Monday - Friday 8:30 AM - 5:00 PM.

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at http://www.uspto.gov/interviewpractice.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Wing Chan can be reached on (571) 272-7493. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/BRENDAN Y HIGA/
Primary Examiner, Art Unit 2441

Notice of References Cited	Application/Control No. 15/836,824	Applicant(s)/Patent Under Reexamination Taaghol et al.	
	Examiner BRENDAN HIGA	Art Unit 2441	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-20190165997-A1	05-2019	SHAIKH; Nazneen	H04L41/0806	1/1
*	B	US-20160224460-A1	08-2016	BRYANT; Andy	G06F8/61	1/1
*	C	US-20150242237-A1	08-2015	Spivak; Vadim	G06F9/5055	718/1
*	D	US-20150160877-A1	06-2015	Larkin; Michael	G06F3/0689	711/114
*	E	US-20120096071-A1	04-2012	Murphey; C. Michael	G06F8/63	709/203
*	F	US-20120066665-A1	03-2012	Kunze; Tobias	G06F8/62	717/124
*	G	US-20110265077-A1	10-2011	COLLISON; Derek	G06F8/71	717/172
	H					
	I					
	J					
	K					
	L					
	M					

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0270

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	BRENDAN HIGA	2441

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	06/16/2019	01/17/2020							
	1	✓	✓							
	2	✓	✓							
	3	✓	-							
	4	✓	-							
	5	✓	✓							
	6	✓	-							
	7	✓	✓							
	8	✓	-							
	9	✓	-							
	10	✓	-							
	11	✓	✓							
	12	✓	✓							
	13	✓	✓							
	14	✓	✓							
	15	✓	✓							
	16	✓	✓							
	17	✓	✓							
	18	✓	✓							

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0271

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	Art Unit
	BRENDAN HIGA	2441

CPC - Searched*		
Symbol	Date	Examiner
H04L 9/08-0819, 32-3212; 29/08135, 0845-08504; 41/08-082, 0866, 0876-0893; 45/563, 64; 67/10,1074-1087	01/17/2020	BH
G06F 8/60-61, 65, 70-71, 76; 9/44006, 45533-4856;	01/17/2020	BH

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	Date	Examiner
see EAST search history	01/17/2020	BH
google scholar search: "distributed software defined network" " distributed software defined networking"; Openflow OR SDN controller "application store" OR "app store";	01/17/2020	BH

/R.K./ Examiner, Art Unit 2441	

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0272

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	Art Unit
	BRENDAN HIGA	2441

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**
H04L	9/08-0819, 32-3212; 29/08135, 0845-08504; 41/08-082, 0866, 0876-0893; 45/563, 64; 67/10,1074-1087	01/17/2020	BH
G06F	8/60-61, 65, 70-71, 76; 9/44006, 45533-4856;	01/17/2020	BH

/R.K./ Examiner, Art Unit 2441	

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0273

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	42956	(VM cloud server) adj25 (application web$1app app) same (open$1flow software adj3 newtork SDN orchestrat$3 controller)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:29
L2	8152	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj15 (application web$1app app) same L1	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:29
L3	330	(version$3 updat$3 recent$3 up$1grad$3) adj15 (application web$1app app) same L2	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:29
L4	222	L3 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:29
L5	109053	(VM cloud server) adj25 (application web$1app app) same (virtual$5 hyper$2visor open$1flow $3SDN software adj2 network$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:57
L6	36934	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj25 (application web$1app app) adj10 (download$3 installation install$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:57
L7	797	L5 same L6	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:57
L8	171	L7 and L1	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:57
L9	153	L8 NOT 4	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 20:57
L10	4770934	(container tenant mult$1tenant sandbox$4 virtual adj2 machine)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT;	OR	ON	2020/01/16 21:04

			IBM_TDB			
L11	8709	(VM cloud server) adj25 (application web$1app app) same (open$1flow $3SDN software adj2 network$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:04
L14	67820	(VM cloud server) adj25 (application web$1app app cloud$1app) same 10	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:06
L15	7226	14 and (open$1flow $3SDN software adj2 defin$3 adj2 network$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:06
L17	435	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj15 (application web$1app app) same 14 and 11	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:07
L18	239	17 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:08
L20	25	(portal merchant market$1place) adj15 (application web$1app cloud$1app app) same 14 and 11	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:10
L21	18	20 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 21:10
S2	3	((("Ramanna") near3 ("Vivek"))).INV.	US-PGPUB; USPAT; USOCR	OR	ON	2020/01/16 16:39
S3	1467	distribut$3 adj2 (Software adj2 defin$3 adj2 Network (SDN)) d$1sdn	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 16:40
S4	109053	(VM cloud server) adj25 (application web$1app app) same (virtual$5 hyper$2visor open$1flow $3SDN software adj2 network$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 16:59
S5	36934	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj25 (application web$1app app) adj10 (download$3 installation install$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 16:59
S6	7	S3 and S4 and S5	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:00

S7	797	S4 same S5	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:03
S8	576	S7 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:04
S9	21048	(version$3 updat$3 recent$3 up$1grad$3) adj15 (application web$1app app) adj25 (download$3 installation install$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:25
S10	137975	(version$3 updat$3 recent$3 up$1grad$3) adj15 (authenticat$3 password authentic key check$1sum authoriz$6)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:26
S11	13	S8 and S9 and S10	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:26
S12	342	S4 same S9	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:28
S13	29	S7 and S12	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:28
S14	15	S13 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:28
S16	42956	(VM cloud server) adj25 (application web$1app app) same (open$1flow software adj3 newtork SDN orchestrat$3 controller)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:34
S17	286	S5 same S16	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:34
S18	8152	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj15 (application web$1app app) same S16	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:35
S19	171	S7 and S16	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:35

S20	146	S19 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:36
S21	330	(version$3 updat$3 recent$3 up$1grad$3) adj15 (application web$1app app) same S18	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/01/16 17:36

EAST Search History (Interference)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S23	5560	(VM cloud server) adj25 (application web$1app app) same (virtual$5 hyper$2visor open$1flow $3SDN software adj2 network$3).clm.	US-PGPUB; USPAT	OR	ON	2020/01/16 17:17
S24	2383	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj25 (application web$1app app) adj10 (download$3 installation install$3).clm.	US-PGPUB; USPAT	OR	ON	2020/01/16 17:17
S25	80	S23 and S24	US-PGPUB; USPAT	OR	ON	2020/01/16 17:17
S26	70	S25 and (@pd @ptad @prad @ad @rlad) <"20130613"	US-PGPUB; USPAT	OR	ON	2020/01/16 17:17
S27	16953	(version$3 updat$3 recent$3 up$1grad$3) adj15 (authenticat$3 password authentic key check$1sum authoriz$6).clm.	US-PGPUB; USPAT	OR	ON	2020/01/16 17:27
S28	0	S26 and S27	US-PGPUB; USPAT	OR	ON	2020/01/16 17:27

1/16/2020 9:12:43 PM
C:\Users\bhiga\Documents\EAST\Workspaces\15836824_1.wsp

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0277

PTO/AIA/26 (04-14)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT	Docket Number (Optional) TAAG001CONT1

In re Application of: Taaghol

Application No.: 15/836,824

Filed: 12/9/2017

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

The applicant, Big Data Federation, Inc., owner of 100 percent interest in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of **prior patent** No. 9843624 as the term of said **prior patent** is presently shortened by any terminal disclaimer. The applicant hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the **prior patent** are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the applicant does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the **prior patent**, "as the term of said **prior patent** is presently shortened by any terminal disclaimer," in the event that said **prior patent** later:

expires for failure to pay a maintenance fee;
is held unenforceable;
is found invalid by a court of competent jurisdiction;
is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
has all claims canceled by a reexamination certificate;
is reissued; or
is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

Check either box 1 or 2 below, if appropriate.

1. ☐ The undersigned is the applicant. If the applicant is an assignee, the undersigned is authorized to act on behalf of the assignee.

I hereby acknowledge that any willful false statements made are punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

2. ☑ The undersigned is an attorney or agent of record. Reg. No. 40,394

/Christopher J. Horgan/ — Signature
1/26/2020 — Date

Christopher J. Horgan — Typed or printed name

Attorney for Applicant — Title
408-687-8306 — Telephone Number

☑ Terminal disclaimer fee under 37 CFR 1.20(d) included.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

This collection of information is required by 37 CFR 1.321. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	15/836,824
Filing Date	12/9/2017
First Named Inventor	Taaghol
Title	Distributed Software Defined Networking
Art Unit	2441
Examiner Name	Higa
Attorney Docket Number	TAAG001CONT1

SIGNATURE of Applicant or Patent Practitioner

Signature	/Christopher J. Horgan/	Date (Optional)	January 27, 2020
Name	Christopher J. Horgan	Registration Number	40,394
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82B (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date
15/836,824	12/9/2017

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

☑ I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above: 112449

OR

☐ I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

☑ The address associated with the above-mentioned Customer Number

OR

☐ The address associated with Customer Number:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

Big Data Federation, Inc.

☐ Inventor or Joint Inventor (title not required below)

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)

☑ Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE of Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature		Date (Optional)	1/26/2020
Name	Pouya Taaghol		
Title	CEO, Big Data Federation, Inc.		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Taaghol
Serial No.: 15/836,824
Filed: 12/09/2017
Attorney Docket No.: TAAG001CONT1

Examiner: Higa, Brendan
Art Unit: 2441

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

Mail Stop Amendment Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450
Dear Sir:

REQUEST UNDER 37 CFR 1.46(c)

Applicants offer the following request to change the Applicant to Big Data Federation, Inc. Applicant is submitting the following:

1) A Supplemental ADS with the Applicant corrected.
2) A new Terminal Disclaimer.
3) A Statement Under 37 CFR 3.73(c).
4) A new Power of Attorney.

No new matter has been added.

Respectfully submitted,
ROARK IP
By: /Christopher J. Horgan/

Christopher J. Horgan
Registration No. 40,394
1438 Dahlia Loop
San Jose, CA 95126
Telephone: (408) 687-8306

Date: January 26, 2020

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0281

Electronic Acknowledgement Receipt

EFS ID:	38400092
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	27-JAN-2020
Filing Date:	09-DEC-2017
Time Stamp:	00:29:30
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Application Data Sheet	TAAG001CONT1_NEW_ADS.pdf	845487 d6bd28beb81a65018ad2d6628f1a7210d5473d4a	no	6

Warnings:

Information:					
This is not an USPTO supplied ADS fillable form					
2	Assignee showing of ownership per 37 CFR 3.73	TAAG001CONT1_CERT.pdf	89497 bd1f817d107762edae7eaf65eab58e5915612473	no	2
Warnings:					
Information:					
3	Terminal Disclaimer Filed	TAAG001CONT1_NEW_TD.pdf	98136 a1983b1b36e095f627a7a65cb08b45871a51816d	no	1
Warnings:					
Information:					
4	Power of Attorney	TAAG001CONT1_POA__.pdf	198703 557b94f72e6d1b0e9f5e1b15b2b4101814ddda7f0	no	2
Warnings:					
Information:					
5	Transmittal Letter	TAAG001CONT1_REQUEST___.pdf	55322 87bfb03e4dd222b52520a3b1aefc47f1ed2b0057	no	1
Warnings:					
Information:					
			Total Files Size (in bytes):	1287145	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fail under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Pouya		Taaghol	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	1260 Trestlewood Lane		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95138	Country	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Vivek		Ramanna	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	549 Manhattan Place		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95136	Country	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

☐ An Address is being provided for the correspondence Information of this application.			
Customer Number	112449		
Email Address	chris.horgan@roarkip.com	Add Email	Remove Email

Application Information:

Title of the Invention	Distributed Software Defined Networking		
Attorney Docket Number	TAAG001CONT1	**Small Entity Status Claimed**	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	5	**Suggested Figure for Publication (if any)**	1

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☒ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	112449		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	Continuation of	14295331	2014-06-04
Prior Application Status	Expired		Remove

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
14295331	non provisional of	61634807	2013-06-13

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)[i] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Remove

Application Number	Country[i]	Filing Date (YYYY-MM-DD)	Access Code[j] (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date o f filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

◯ Assignee	◯ Legal Representative under 35 U.S.C. 117	◯ Joint Inventor
◯ Person to whom the inventor is obligated to assign.	◯ Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. ☐

Prefix	Given Name	Middle Name	Family Name	Suffix

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Mailing Address Information For Applicant:			
Address 1			
Address 2			
City		State/Province	
Country[i]		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Applicant Data may be generated within this form by selecting the Add button.

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

If the Assignee is an Organization check here. ☐

Prefix	Given Name	Middle Name	Family Name	Suffix

Mailing Address Information For Non-Applicant Assignee:			
Address 1			
Address 2			
City		State/Province	
Country[i]		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button.

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Signature:

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications.					
Signature	/Christopher J. Horgan/			Date (YYYY-MM-DD)	2017-12-08
First Name	Christopher	Last Name	Horgan	Registration Number	40394
Additional Signature may be generated within this form by selecting the Add button.					

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

PTO/AIA/96 (08-12)
Approved for use through 01/31/2013. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(c)

Applicant/Patent Owner: Big Data Federation, Inc.

Application No./Patent No.: 15/836,824 Filed/Issue Date: 12/9/2017

Titled: Distributed Software Defined Networking

Big Data Federation, Inc., a corporation

(Name of Assignee) (Type of Assignee, e.g., corporation, partnership, university, government agency, etc.)

states that, for the patent application/patent identified above, it is (choose **one** of options 1, 2, 3 or 4 below):

1. ☑ The assignee of the entire right, title, and interest.

2. ☐ An assignee of less than the entire right, title, and interest (check applicable box):

   ☐ The extent (by percentage) of its ownership interest is _____%. Additional Statement(s) by the owners holding the balance of the interest must be submitted to account for 100% of the ownership interest.

   ☐ There are unspecified percentages of ownership. The other parties, including inventors, who together own the entire right, title and interest are:

   Additional Statement(s) by the owner(s) holding the balance of the interest must be submitted to account for the entire right, title, and interest.

3. ☐ The assignee of an undivided interest in the entirety (a complete assignment from one of the joint inventors was made). The other parties, including inventors, who together own the entire right, title, and interest are:

   Additional Statement(s) by the owner(s) holding the balance of the interest must be submitted to account for the entire right, title, and interest.

4. ☐ The recipient, via a court proceeding or the like (*e.g.*, bankruptcy, probate), of an undivided interest in the entirety (a complete transfer of ownership interest was made). The certified document(s) showing the transfer is attached.

The interest identified in option 1, 2 or 3 above (not option 4) is evidenced by either (choose **one** of options A or B below):

A. ☐ An assignment from the inventor(s) of the patent application/patent identified above. The assignment was recorded in the United States Patent and Trademark Office at Reel _____, Frame _____, or for which a copy thereof is attached.

B. ☑ A chain of title from the inventor(s), of the patent application/patent identified above, to the current assignee as follows:

1. From: Pouya Taaghol To: Big Data Federation, Inc.

   The document was recorded in the United States Patent and Trademark Office at
   Reel 050963, Frame 0054, or for which a copy thereof is attached.

2. From: Vivek Ramanna To: Pouya Taaghol

   The document was recorded in the United States Patent and Trademark Office at
   Reel 035569, Frame 0334, or for which a copy thereof is attached.

This collection of information is required by 37 CFR 3.73(b). The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/96 (08-12)
Approved for use through 01/31/2013. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(c)

3. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

4. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

5. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

6. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

☐ Additional documents in the chain of title are listed on a supplemental sheet(s).

☑ As required by 37 CFR 3.73(c)(1)(i), the documentary evidence of the chain of title from the original owner to the assignee was, or concurrently is being, submitted for recordation pursuant to 37 CFR 3.11.

[NOTE: A separate copy (i.e., a true copy of the original assignment document(s)) must be submitted to Assignment Division in accordance with 37 CFR Part 3, to record the assignment in the records of the USPTO. See MPEP 302.08]

The undersigned (whose title is supplied below) is authorized to act on behalf of the assignee.

/Christopher J. Horgan/	January 26, 2020
Signature	Date
Christopher J. Horgan	40,394
Printed or Typed Name	Title or Registration Number

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1

CONFIRMATION NO. 9801

IMPROPER CPOA LETTER

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

OC000000114310573

Date Mailed: 01/29/2020

NOTICE REGARDING POWER OF ATTORNEY

This is in response to the power of attorney filed 01/27/2020. The power of attorney in this application is not accepted for the reason(s) listed below:

- The power of attorney has not been accepted because the party who is giving power has not been identified. Power of attorney may only be signed by the applicant for patent (37 CFR 1.42) or the patent owner. A party who is not the applicant must become the applicant in accordance with 37 CFR 1.46(c) and appoint any power of attorney in compliance with 37 CFR 3.71 and 3.73. For a reissue application, reexamination proceeding, or supplemental examination proceeding, a patent owner who was not the applicant under 37 CFR 1.46 must appoint any power of attorney in compliance with 37 CFR 3.71 and 3.73. See 37 CFR 1.32(b)(4).

/sstephanos/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126
UNITED STATES

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1

CONFIRMATION NO. 9801

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

IMPROPER CFR REQUEST

OC000000114311186

Date Mailed: 01/29/2020

RESPONSE TO REQUEST FOR CORRECTED FILING RECEIPT

Power of Attorney, Claims, Fees, System Limitations, and Miscellaneous

In response to your request for a corrected Filing Receipt, the Office is unable to comply with your request because:

- The ADS submitted on 01/27/2020 was not properly marked up to show the desired changes. For information being changed relative to the information already of record, additions must be shown with underlining, and deletions must be shown with strike-through or brackets. See 37 CFR 1.76(c)(2)

In order to make changes to the information of record, an ADS must be properly signed and properly marked up relative to the current information of record.

Proper signature: The ADS must be signed with a handwritten signature or proper S-signature by:

- A patent practitioner, with the practitioner's registration number accompanying the signature (e.g., immediately below or adjacent to the signature), or
- The applicant, if the applicant is an individual other than the inventor(s) and no power of attorney has been appointed, or
- All of the inventors, if no other applicant has been established and no power of attorney has been appointed.

A proper S-signature consists of only letters and/or Arabic numerals, with appropriate spaces and commas, periods, apostrophes, or hyphens for punctuation contained between a first single forward slash mark before, and a second single forward slash mark after, the S-signature.

Proper markings: The ADS must identify the changes being made with underlining for insertions and strike-through or brackets for text removed. No other markings or indications are acceptable. Where an ADS providing corrected or updated information does not contain all of the sections of the ADS, the entire section in which changes are being made must be included in the ADS. Information of record can generally be found on the latest filing receipt.

Questions about the contents of this notice and the requirements it sets forth should be directed to the Office of Data Management, Application Assistance Unit, at **(571) 272-4000** or **(571) 272-4200** or **1-888-786-0101**.

/sstephanos/

Application Number	Application/Control No.	Applicant(s)/Patent under Reexamination	
* 15/836,824 *	15/836,824	Taaghol et al.	
	Examiner	**Art Unit**	
	HIGA, BRENDAN Y	2441	

Document Code - DISQ	**Internal Document - DO NOT MAIL**

TERMINAL DISCLAIMER	☐ APPROVED	☑ DISAPPROVED
Date Filed: **27 January 2020**	**This patent is subject to a Terminal Disclaimer**	

Approved/Disapproved by:

/FELICIA D ROBERTS/

Technology Center: OPLC

Telephone: (571)272-0550

See Continuation Sheet

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0294

Continuation of Approved/Disapproved Comments: The terminal disclaimer identifies a party who is not the applicant (only for applications filed on or after September 16, 2012; See FP 14.26.10):

For cases filed on/after 9/16/12, 37 CFR 1.321 specifies that the applicant can disclaim, and the terminal disclaimer must specify the extent of the applicant's ownership.

To remedy this:

A request under 37 CFR 1.46(c) to change the applicant needs to be filed, which is (1) a request, signed by a 1.33(b) party, (2) a corrected ADS (37 CFR 1.76(c)) that identifies the "new" applicant in the applicant information, and is underlined since it is new, and (3) a 3.73(c) statement showing chain of title to the new applicant. Along with the § 1.46(c) request we need a POA that gives power to the attorney who is signing the TD, along with another copy of the TD, unless they file a TD that is signed by the applicant.

Please make corrections as suggested above and also resubmit the TD. (No new fee required)

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	15/836,824
Filing Date	12/9/2017
First Named Inventor	Taaghol
Title	Distributed Software Defined Networking
Art Unit	2441
Examiner Name	Higa
Attorney Docket Number	TAAG001CONT1

SIGNATURE of Applicant or Patent Practitioner

Signature	/Christopher J. Horgan/	Date (Optional)	March 2, 2020
Name	Christopher J. Horgan	Registration Number	40,394
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82B (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date
15/836,824	12/9/2017

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

☑ I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above: 112449

OR

☐ I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

☑ The address associated with the above-mentioned Customer Number

OR

☐ The address associated with Customer Number:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

Big Data Federation, Inc.

☐ Inventor or Joint Inventor (title not required below)

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)

☑ Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE of Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature		Date (Optional)	1/26/2020
Name	Pouya Taaghol		
Title	CEO, Big Data Federation, Inc.		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/26 (04-14)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT	Docket Number (Optional) TAAG001CONT1

In re Application of: Taaghol

Application No.: 15/836,824

Filed: 12/9/2017

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

The applicant, Big Data Federation, Inc., owner of 100 percent interest in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of **prior patent** No. 9843624 as the term of said **prior patent** is presently shortened by any terminal disclaimer. The applicant hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the **prior patent** are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the applicant does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the **prior patent**, "as the term of said **prior patent** is presently shortened by any terminal disclaimer," in the event that said **prior patent** later:

expires for failure to pay a maintenance fee;
is held unenforceable;
is found invalid by a court of competent jurisdiction;
is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
has all claims canceled by a reexamination certificate;
is reissued; or
is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

Check either box 1 or 2 below, if appropriate.

1. ☐ The undersigned is the applicant. If the applicant is an assignee, the undersigned is authorized to act on behalf of the assignee.

I hereby acknowledge that any willful false statements made are punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

2. ☑ The undersigned is an attorney or agent of record. Reg. No. 40,394

/Christopher J. Horgan/ — Signature

3/2/2020 — Date

Christopher J. Horgan — Typed or printed name

Attorney for Applicant — Title

408-687-8306 — Telephone Number

☑ Terminal disclaimer fee under 37 CFR 1.20(d) included.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

This collection of information is required by 37 CFR 1.321. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Taaghol
Serial No.: 15/836,824
Filed: 12/09/2017
Attorney Docket No.: TAAG001CONT1

Examiner: Higa, Brendan
Art Unit: 2441

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

Mail Stop Amendment Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450
Dear Sir:

RENEWED REQUEST UNDER 37 CFR 1.46(c)

Applicants offer the following request to change the Applicant to Big Data Federation, Inc. Applicant is submitting the following:

1) A Supplemental ADS with the Applicant corrected.
2) A new Terminal Disclaimer.
3) A new Power of Attorney.

No new matter has been added.

Respectfully submitted,
ROARK IP
By: /Christopher J. Horgan/

Christopher J. Horgan
Registration No. 40,394
1438 Dahlia Loop
San Jose, CA 95126
Telephone: (408) 687-8306

Date: March 3, 2020

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0299

Electronic Acknowledgement Receipt

EFS ID:	38754923
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	03-MAR-2020
Filing Date:	09-DEC-2017
Time Stamp:	14:28:16
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Application Data Sheet	TAAG001CONT1_SUPP_ADS__.pdf	1532613 d0d26a6d442461f3f32d52a7145072f5de6a1dd0	no	5

Warnings:

Information:					
This is not an USPTO supplied ADS fillable form					
2	Power of Attorney	TAAG001CONT1_NEW_POA__.pdf	197572 42e39b5d639d291bc5bd9f4c18012f5ae8bbc901	no	2
Warnings:					
Information:					
3	Application Data Sheet	TAAG001CONT1_NEW_TD__.pdf	84597 87b94acf88601c75b20f13e83a792d8505aaade0	no	1
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
4	Req to correct/update/change appl name until pay issue fee	TAAG001CONT1_REQUEST___.pdf	35797 f3eecfda712be8f6259b9621f15288c0f9c55f70	no	1
Warnings:					
Information:					
		Total Files Size (in bytes):	1850579		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Supplemental

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Pouya		Taaghol	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	1260 Trestlewood Lane		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95138	Country	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Vivek		Ramanna	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	549 Manhattan Place		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95126	Country	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

☐ An Address is being provided for the correspondence Information of this application.			
Customer Number	112449		
Email Address	chris.horgan@roarkip.com	Add Email	Remove Email

Application Information:

Title of the Invention	Distributed Software Defined Networking		
Attorney Docket Number	TAAG001CONT1	**Small Entity Status Claimed**	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	5	**Suggested Figure for Publication (if any)**	1

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☒ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	112449		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	Continuation of	14295331	2014-06-04
Prior Application Status	Pending		Remove

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
14295331	non provisional of	61834807	2013-06-13

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)[1] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Remove

Application Number	Country[1]	Filing Date (YYYY-MM-DD)	Access Code[1] (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date o f filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

(•) Assignee	( ) Legal Representative under 35 U.S.C. 117	( ) Joint Inventor
( ) Person to whom the inventor is obligated to assign.	( ) Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :	
If the Applicant is an Organization check here.	☒
Organization Name	Big Data Federation, Inc.

Mailing Address Information For Applicant:

Address 1	990 Linden Drive, #220		
Address 2			
City	Santa Clara	**State/Province**	CA
Country	US	Postal Code	95050
Phone Number	408-623-8529	Fax Number	

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Email Address	pouya@bigdatafed.com

Additional Applicant Data may be generated within this form by selecting the Add button.

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

If the Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Non-Applicant Assignee:

Address 1			
Address 2			
City		**State/Province**	
Country		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button.

Signature:

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Signature	/Christopher J. Horgan/			Date (YYYY-MM-DD)	2020-03-02
First Name	Christopher	Last Name	Horgan	Registration Number	40394

Additional Signature may be generated within this form by selecting the Add button.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1

CONFIRMATION NO. 9801

IMPROPER CPOA LETTER

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

OC000000115210521

Date Mailed: 03/05/2020

NOTICE REGARDING POWER OF ATTORNEY

This is in response to the power of attorney filed 03/03/2020. The power of attorney in this application is not accepted for the reason(s) listed below:

- The power of attorney has not been accepted because the party who is giving power has not been identified. Power of attorney may only be signed by the applicant for patent (37 CFR 1.42) or the patent owner. A party who is not the applicant must become the applicant in accordance with 37 CFR 1.46(c) and appoint any power of attorney in compliance with 37 CFR 3.71 and 3.73. For a reissue application, reexamination proceeding, or supplemental examination proceeding, a patent owner who was not the applicant under 37 CFR 1.46 must appoint any power of attorney in compliance with 37 CFR 3.71 and 3.73. See 37 CFR 1.32(b)(4).

The power of attorney filed 03/03/2020 has not been accepted because the power of attorney must be signed by the applicant for patent. See 37 CFR 1.32(b)(4).

- The person or entity attempting the change to power of attorney is not the applicant of record in the application. Any request to change the applicant once the applicant has been specified must include (1) an application data sheet (ADS) specifying the new applicant in the Applicant Information section, and (2) a statement under 37 CFR 3.73(c) (USPTO Form PTO/AIA/96 or an equivalent) to show chain of title to the new applicant. The ADS must contain markings to show the information that is being changed, with underlining for insertions and strike-through or brackets for text removed. See 37 CFR 1.76(c)(2).

Because the request to change or update the applicant cannot be accepted, the power of attorney is not properly signed by the applicant and cannot be accepted. The change to applicant must be acceptable before the new applicant can appoint power of attorney.

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126
UNITED STATES

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0307

/dtdinh/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0308

PTO/AIA/26 (04-14)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT

Docket Number (Optional): TAAG001CONT1

In re Application of: Taaghol

Application No.: 15/836,824

Filed: 12/9/2017

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

The applicant, Big Data Federation, Inc., owner of 100 percent interest in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of **prior patent** No. 9843624 as the term of said **prior patent** is presently shortened by any terminal disclaimer. The applicant hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the **prior patent** are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the applicant does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the **prior patent**, "as the term of said **prior patent** is presently shortened by any terminal disclaimer," in the event that said **prior patent** later:

expires for failure to pay a maintenance fee;
is held unenforceable;
is found invalid by a court of competent jurisdiction;
is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
has all claims canceled by a reexamination certificate;
is reissued; or
is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

Check either box 1 or 2 below, if appropriate.

1. ☐ The undersigned is the applicant. If the applicant is an assignee, the undersigned is authorized to act on behalf of the assignee.

I hereby acknowledge that any willful false statements made are punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

2. ☑ The undersigned is an attorney or agent of record. Reg. No. 40,394

/Christopher J. Horgan/ — Signature

4/17/2020 — Date

Christopher J. Horgan — Typed or printed name

Attorney for Applicant — Title

408-687-8306 — Telephone Number

☑ Terminal disclaimer fee under 37 CFR 1.20(d) included.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

This collection of information is required by 37 CFR 1.321. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/96 (08-12)
Approved for use through 01/31/2013. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(c)

Applicant/Patent Owner: Big Data Federation, Inc.

Application No./Patent No.: 15/836,824 Filed/Issue Date: 12/9/2017

Titled: Distributed Software Defined Networking

Big Data Federation, Inc., a corporation

(Name of Assignee) (Type of Assignee, e.g., corporation, partnership, university, government agency, etc.)

states that, for the patent application/patent identified above, it is (choose **one** of options 1, 2, 3 or 4 below):

1. ☑ The assignee of the entire right, title, and interest.

2. ☐ An assignee of less than the entire right, title, and interest (check applicable box):

   ☐ The extent (by percentage) of its ownership interest is _____%. Additional Statement(s) by the owners holding the balance of the interest must be submitted to account for 100% of the ownership interest.

   ☐ There are unspecified percentages of ownership. The other parties, including inventors, who together own the entire right, title and interest are:

   Additional Statement(s) by the owner(s) holding the balance of the interest must be submitted to account for the entire right, title, and interest.

3. ☐ The assignee of an undivided interest in the entirety (a complete assignment from one of the joint inventors was made). The other parties, including inventors, who together own the entire right, title, and interest are:

   Additional Statement(s) by the owner(s) holding the balance of the interest must be submitted to account for the entire right, title, and interest.

4. ☐ The recipient, via a court proceeding or the like (*e.g.*, bankruptcy, probate), of an undivided interest in the entirety (a complete transfer of ownership interest was made). The certified document(s) showing the transfer is attached.

The interest identified in option 1, 2 or 3 above (not option 4) is evidenced by either (choose **one** of options A or B below):

A. ☐ An assignment from the inventor(s) of the patent application/patent identified above. The assignment was recorded in the United States Patent and Trademark Office at Reel _____, Frame _____, or for which a copy thereof is attached.

B. ☑ A chain of title from the inventor(s), of the patent application/patent identified above, to the current assignee as follows:

1. From: Pouya Taaghol To: Big Data Federation, Inc.

   The document was recorded in the United States Patent and Trademark Office at
   Reel 050963, Frame 0054, or for which a copy thereof is attached.

2. From: Vivek Ramanna To: Pouya Taaghol

   The document was recorded in the United States Patent and Trademark Office at
   Reel 035569, Frame 0334, or for which a copy thereof is attached.

This collection of information is required by 37 CFR 3.73(b). The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/96 (08-12)
Approved for use through 01/31/2013. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(c)

3. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

4. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

5. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

6. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

☐ Additional documents in the chain of title are listed on a supplemental sheet(s).

☑ As required by 37 CFR 3.73(c)(1)(i), the documentary evidence of the chain of title from the original owner to the assignee was, or concurrently is being, submitted for recordation pursuant to 37 CFR 3.11.

[NOTE: A separate copy (i.e., a true copy of the original assignment document(s)) must be submitted to Assignment Division in accordance with 37 CFR Part 3, to record the assignment in the records of the USPTO. See MPEP 302.08]

The undersigned (whose title is supplied below) is authorized to act on behalf of the assignee.

/Christopher J. Horgan/	April 17, 2020
Signature	Date
Christopher J. Horgan	40,394
Printed or Typed Name	Title or Registration Number

Electronic Patent Application Fee Transmittal

Application Number:	15836824
Filing Date:	09-Dec-2017
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Filer:	Christopher James Horgan
Attorney Docket Number:	TAAG001CONT1

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension - 1 month with $0 paid	2251	1	100	100
Miscellaneous:				
		Total in USD ($)		**100**

Electronic Acknowledgement Receipt

EFS ID:	39196448
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	17-APR-2020
Filing Date:	09-DEC-2017
Time Stamp:	18:19:21
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$100
RAM confirmation Number	E20204GI23538947
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Req to correct/update/change appl name until pay issue fee	TAAG001CONT1_LETTER__.pdf	18719 60e263869243c23059de8ab3d0406f7a60b88250	no	1
Warnings:					
Information:					
2	Application Data Sheet	TAAG001CONT1_NEW_ADS__.pdf	910981 ad51a89f380b94b1032a10b2b10be0703ebe5485	no	5
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
3	Power of Attorney	TAAG001CONT1_NEW_POA__.pdf	178160 ac00991a694205caa589bbd69d698afff3ddaccc	no	2
Warnings:					
Information:					
4	Terminal Disclaimer Filed	TAAG001CONT1_NEW_TD__.pdf	83265 b06abebe6e175399c4abc37fd67a1fbf3f31d2e4	no	1
Warnings:					
Information:					
5	Assignee showing of ownership per 37 CFR 3.73	TAAG001CONT1_CERT__.pdf	170268 10d0468c0c2ba5f3b07d779d79302eec04bbc268	no	2
Warnings:					
Information:					
6	Fee Worksheet (SB06)	fee-info.pdf	30796 53ce70a95285cd7d81921b81635e12577230c387	no	2
Warnings:					

Information:	
Total Files Size (in bytes):	1392189

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Taaghol
Serial No.: 15/836,824
Filed: 12/09/2017
Attorney Docket No.: TAAG001CONT1

Examiner: Higa, Brendan
Art Unit: 2441

For: DISTRIBUTED SOFTWARE DEFINED NETWORKING

Mail Stop Amendment Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450
Dear Sir:

RENEWED REQUEST UNDER 37 CFR 1.46(c)

Applicants offer the following request to change the Applicant to Big Data Federation, Inc. Applicant is submitting the following:

1) A Supplemental ADS with the Applicant corrected.
2) A new Terminal Disclaimer.
3) A new Power of Attorney.
4) Statement Under 37 CFR 3.73(c).

No new matter has been added.

Respectfully submitted,
ROARK IP
By: /Christopher J. Horgan/

Christopher J. Horgan
Registration No. 40,394
1438 Dahlia Loop
San Jose, CA 95126
Telephone: (408) 687-8306

Date: April 17, 2020

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0317

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Supplemental

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Pouya		Taaghol	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	1260 Trestlewood Lane		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95138	Country	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Vivek		Ramanna	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	San Jose	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	549 Manhattan Place		
Address 2			
City	San Jose	State/Province	CA
Postal Code	95126	Country	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

☐ An Address is being provided for the correspondence Information of this application.			
Customer Number	112449		
Email Address	chris.horgan@roarkip.com	Add Email	Remove Email

Application Information:

Title of the Invention	Distributed Software Defined Networking		
Attorney Docket Number	TAAG001CONT1	**Small Entity Status Claimed**	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	5	**Suggested Figure for Publication (if any)**	1

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☒ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	112449		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	Continuation of	14295331	2014-06-04
Prior Application Status	Pending		Remove

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
14295331	non provisional of	61834807	2013-06-13

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)[1] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Remove

Application Number	Country[1]	Filing Date (YYYY-MM-DD)	Access Code[1] (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date o f filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

(•) Assignee	( ) Legal Representative under 35 U.S.C. 117	( ) Joint Inventor
( ) Person to whom the inventor is obligated to assign.	( ) Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :	
If the Applicant is an Organization check here.	☒
Organization Name	Big Data Federation, Inc.

Mailing Address Information For Applicant:

Address 1	990 Linden Drive, #220		
Address 2			
City	Santa Clara	**State/Province**	CA
Country	US	Postal Code	95050
Phone Number	408-623-8529	Fax Number	

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	TAAG001CONT1
	Application Number	
Title of Invention	Distributed Software Defined Networking	

Email Address	pouya@bigdatafed.com

Additional Applicant Data may be generated within this form by selecting the Add button.

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

If the Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Non-Applicant Assignee:

Address 1			
Address 2			
City		**State/Province**	
Country i		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button.

Signature:

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Signature	/Christopher J. Horgan/			Date (YYYY-MM-DD)	04-17-2020
First Name	Christopher	Last Name	Horgan	Registration Number	40394

Additional Signature may be generated within this form by selecting the Add button.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	15/836,824
Filing Date	12/9/2017
First Named Inventor	Taaghol
Title	Distributed Software Defined Networking
Art Unit	2441
Examiner Name	Higa
Attorney Docket Number	TAAG001CONT1

SIGNATURE of Applicant or Patent Practitioner

Signature	/Christopher J. Horgan/	Date (Optional)	April 17, 2020
Name	Christopher J. Horgan	Registration Number	40,394
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82B (07-13)
Approved for use through 03/31/2021. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date
15/836,824	12/9/2017

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

☑ I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above: 112449

OR

☐ I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

☑ The address associated with the above-mentioned Customer Number

OR

☐ The address associated with Customer Number:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

Big Data Federation, Inc.

☐ Inventor or Joint Inventor (title not required below)

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)

☑ Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE of Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature		Date (Optional)	1/26/2020
Name	Pouya Taaghol		
Title	CEO, Big Data Federation, Inc.		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Application Number	Application/Control No.	Applicant(s)/Patent under Reexamination	
* 15/836,824 *	15/836,824	Taaghol et al.	
	Examiner	**Art Unit**	
	HIGA, BRENDAN Y	2441	

Document Code - DISQ	**Internal Document - DO NOT MAIL**

TERMINAL DISCLAIMER	☑ APPROVED	☐ DISAPPROVED
Date Filed: 17 April 2020	**This patent is subject to a Terminal Disclaimer**	

Approved/Disapproved by:

/FELICIA D ROBERTS/

Technology Center: OPLC

Telephone: (571)272-0550

9843624

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0325

Application Number	Application/Control No.	Applicant(s)/Patent under Reexamination	
* 15/836,824 *	15/836,824	Taaghol et al.	
	Examiner	**Art Unit**	
	HIGA, BRENDAN Y	2441	

Document Code - DISQ	**Internal Document - DO NOT MAIL**

TERMINAL DISCLAIMER	☐ APPROVED	☑ DISAPPROVED
Date Filed: 05 March 2020	**This patent is subject to a Terminal Disclaimer**	

Approved/Disapproved by:

/FELICIA D ROBERTS/

Technology Center: OPLC

Telephone: (571)272-0550

The person who signed the terminal disclaimer is not the applicant, patentee or an attorney or agent of record. 37 CFR 1.321(a) and (b). (See FP 14.26.08)

Please file a POA that gives power to the attorney who is signing the TD, along with another copy of the TD, or file a TD that is signed by the applicant. (No new fee required) Thanks.

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0326

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1

CONFIRMATION NO. 9801

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

POA ACCEPTANCE LETTER

OC000000116399887

Date Mailed: 04/23/2020

NOTICE OF ACCEPTANCE OF POWER OF ATTORNEY

This is in response to the Power of Attorney filed 04/17/2020.

The Power of Attorney in this application is accepted. Correspondence in this application will be mailed to the above address as provided by 37 CFR 1.33.

Questions about the contents of this notice and the requirements it sets forth should be directed to the Office of Data Management, Application Assistance Unit, at **(571) 272-4000** or **(571) 272-4200** or **1-888-786-0101**.

/mmasfaw/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
15/836,824	12/09/2017	2441	730	TAAG001CONT1	18	3

CONFIRMATION NO. 9801

CORRECTED FILING RECEIPT

112449
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

OC000000118399934

Date Mailed: 04/23/2020

Receipt is acknowledged of this non-provisional utility patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF FIRST INVENTOR, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection.

Please verify the accuracy of the data presented on this receipt. If an error is noted on this Filing Receipt, please submit a written request for a corrected Filing Receipt, including a properly marked-up ADS showing the changes with strike-through for deletions and underlining for additions. If you received a "Notice to File Missing Parts" or other Notice requiring a response for this application, please submit any request for correction to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections provided that the request is grantable.

Inventor(s)

Pouya Taaghol, San Jose, CA;
Vivek Ramanna, San Jose, CA;

Applicant(s)

Big Data Federation, Inc., Santa Clara, CA;

Power of Attorney: The patent practitioners associated with Customer Number 112449

Domestic Priority data as claimed by applicant

This application is a CON of 14/295,331 06/04/2014 PAT 9843624
which claims benefit of 61/834,807 06/13/2013

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access Application via Priority Document Exchange: No

Permission to Access Search Results: No

Applicant may provide or rescind an authorization for access using Form PTO/SB/39 or Form PTO/SB/69 as appropriate.

If Required, Foreign Filing License Granted: 01/10/2018

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 15/836,824**

Projected Publication Date: Request for Non-Publication Acknowledged

Non-Publication Request: Yes

Early Publication Request: No
**** SMALL ENTITY ****
Title

DISTRIBUTED SOFTWARE DEFINED NETWORKING

Preliminary Class

709

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0329

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

112449 7590 05/04/2020

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

EXAMINER
HIGA, BRENDAN Y

ART UNIT	PAPER NUMBER
2441	

DATE MAILED: 05/04/2020

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

TITLE OF INVENTION: DISTRIBUTED SOFTWARE DEFINED NETWORKING

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$500	$0.00	$0.00	$500	08/04/2020

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the ENTITY STATUS shown above. If the ENTITY STATUS is shown as SMALL or MICRO, verify whether entitlement to that entity status still applies.

If the ENTITY STATUS is the same as shown above, pay the TOTAL FEE(S) DUE shown above.

If the ENTITY STATUS is changed from that shown above, on PART B - FEE(S) TRANSMITTAL, complete section number 5 titled "Change in Entity Status (from status indicated above)".

For purposes of this notice, small entity fees are 1/2 the amount of undiscounted fees, and micro entity fees are 1/2 the amount of small entity fees.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Maintenance fees are due in utility patents issuing on applications filed on or after Dec. 12, 1980. It is patentee's responsibility to ensure timely payment of maintenance fees when due. More information is available at www.uspto.gov/PatentMaintenanceFees.

PTOL-85 (Rev. 02/11)

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0331

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), by mail or fax, or via EFS-Web.

By mail, send to: Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

By fax, send to: (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

112449 7590 05/04/2020

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being transmitted to the USPTO via EFS-Web or by facsimile to (571) 273-2885, on the date below.

(Typed or printed name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

TITLE OF INVENTION: DISTRIBUTED SOFTWARE DEFINED NETWORKING

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$500	$0.00	$0.00	$500	08/04/2020

EXAMINER	ART UNIT	CLASS-SUBCLASS
HIGA, BRENDAN Y	2441	709-201000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❑ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-09 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1_____
2_____
3_____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document must have been previously recorded, or filed for recordation, as set forth in 37 CFR 3.11 and 37 CFR 3.81(a). Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE (B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ❑ Corporation or other private group entity ❑ Government

4a. Fees submitted: ❑Issue Fee ❑Publication Fee (if required) ❑Advance Order - # of Copies _____

4b. Method of Payment: *(Please first reapply any previously paid fee shown above)*

❑ Electronic Payment via EFS-Web ❑ Enclosed check ❑ Non-electronic payment by credit card (Attach form PTO-2038)

❑ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment to Deposit Account No. _____

5. **Change in Entity Status** (from status indicated above)

❑ Applicant certifying micro entity status. See 37 CFR 1.29

❑ Applicant asserting small entity status. See 37 CFR 1.27

❑ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature _____ Date _____

Typed or printed name _____ Registration No. _____

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0332

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

112449 7590 05/04/2020

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

EXAMINER	
HIGA, BRENDAN Y	
ART UNIT	PAPER NUMBER
2441	

DATE MAILED: 05/04/2020

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(Applications filed on or after May 29, 2000)

The Office has discontinued providing a Patent Term Adjustment (PTA) calculation with the Notice of Allowance.

Section 1(h)(2) of the AIA Technical Corrections Act amended 35 U.S.C. 154(b)(3)(B)(i) to eliminate the requirement that the Office provide a patent term adjustment determination with the notice of allowance. See Revisions to Patent Term Adjustment, 78 Fed. Reg. 19416, 19417 (Apr. 1, 2013). Therefore, the Office is no longer providing an initial patent term adjustment determination with the notice of allowance. The Office will continue to provide a patent term adjustment determination with the Issue Notification Letter that is mailed to applicant approximately three weeks prior to the issue date of the patent, and will include the patent term adjustment on the patent. Any request for reconsideration of the patent term adjustment determination (or reinstatement of patent term adjustment) should follow the process outlined in 37 CFR 1.705.

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

PTOL-85 (Rev. 02/11)

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0333

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Notice of Allowability	Application No. 15/836,824	Applicant(s) Taaghol et al.	
	Examiner BRENDAN HIGA	Art Unit 2441	AIA (FITF) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☑ This communication is responsive to Applicant's Terminal Disclaimer filed on 04/20/2020.
   ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

3. ☑ The allowed claim(s) is/are 1-2,5,7 and 11-18 . As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information , please see **http://www.uspto.gov/patents/init_events/pph/index.jsp** or send an inquiry to **PPHfeedback@uspto.gov.**

4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   Certified copies:

   a) ☐ All b) ☐ Some *c) ☐ None of the:

   1. ☐ Certified copies of the priority documents have been received.
   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
   3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____ .

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____ .

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☐ Notice of References Cited (PTO-892)
2. ☐ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date _____.
3. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material _____.
4. ☐ Interview Summary (PTO-413), Paper No./Mail Date. _____.
5. ☐ Examiner's Amendment/Comment
6. ☐ Examiner's Statement of Reasons for Allowance
7. ☐ Other _____.

/BRENDAN Y HIGA/
Primary Examiner, Art Unit 2441

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	BRENDAN HIGA	2441

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

CLAIMS

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☑ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	06/16/2019	01/17/2020	04/26/2020						
1	1	✓	✓	=						
2	2	✓	✓	=						
	3	✓	-	-						
	4	✓	-	-						
3	5	✓	✓	=						
	6	✓	-	-						
4	7	✓	✓	=						
	8	✓	-	-						
	9	✓	-	-						
	10	✓	-	-						
5	11	✓	✓	=						
6	12	✓	✓	=						
7	13	✓	✓	=						
8	14	✓	✓	=						
9	15	✓	✓	=						
10	16	✓	✓	=						
11	17	✓	✓	=						
12	18	✓	✓	=						

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0336

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	Art Unit
	BRENDAN HIGA	2441

CPC - Searched*		
Symbol	Date	Examiner
H04L 9/08-0819, 32-3213; 29/08135, 0845-08459; 41/08-082, 0866-0893; 45/563, 64; 63/0218; 67/10,1074-1087	04/27/2020	BH
G06F 8/60-61, 65, 70-71, 76; 9/4406, 45533-45558; 21/44-53' 2009/ 45562-45595;	04/27/2020	BH

CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	Date	Examiner
see EAST search history	04/27/2020	BH
google scholar search: "distributed software defined network" " distributed software defined networking"; Openflow OR SDN controller "application store" OR "app store";	01/17/2020	BH

/R.K./ Examiner, Art Unit 2441	

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0337

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	Art Unit
	BRENDAN HIGA	2441

Interference Search			
US Class/CPC Symbol	**US Subclass/CPC Group**	**Date**	**Examiner**
H04L	9/08-0819, 32-3213; 29/08135, 0845-08459; 41/08-082, 0866-0893; 45/563, 64; 63/0218; 67/10,1074-1087	04/27/2020	BH
G06F	8/60-61, 65, 70-71, 76; 9/4406, 45533-45558; 21/44-53' 2009/45562-45595;	04/27/2020	BH

/R.K./ Examiner, Art Unit 2441	

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0338

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	BRENDAN HIGA	2441

CPC						
Symbol					**Type**	**Version**
H04L	/	67	/	10	F	2013-01-01
H04L	/	41	/	042	I	2013-01-01
G06F	/	9	/	4406	I	2013-01-01
G06F	/	21	/	53	I	2013-01-01
H04L	/	29	/	08135	I	2013-01-01
H04L	/	41	/	0893	I	2013-01-01
H04L	/	29	/	08459	I	2013-01-01
H04L	/	9	/	0822	I	2013-01-01
H04L	/	45	/	64	I	2013-01-01
H04L	/	45	/	563	I	2013-01-01
H04L	/	67	/	1076	I	2013-01-01
G06F	/	8	/	60	A	2013-01-01
H04L	/	63	/	0218	A	2013-01-01
G06F	/	21	/	44	A	2013-01-01
H04W	/	4	/	00	A	2013-01-01
G06F	/	21	/	57	A	2013-01-01
G06F	/	2009	/	45575	A	2013-01-01
G06F	/	8	/	70	A	2013-01-01

CPC Combination Sets							
Symbol					**Type**	**Set**	**Ranking**	**Version**
	/		/					

NONE		Total Claims Allowed:	
(Assistant Examiner)	(Date)	12	
/BRENDAN Y HIGA/ Primary Examiner, Art Unit 2441	27 April 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	3

U.S. Patent and Trademark Office — Part of Paper No.: 20200426

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0339

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	BRENDAN HIGA	2441

INTERNATIONAL CLASSIFICATION				
CLAIMED				
H04L29/08	/	29	/	08
H04L12/24	/	12	/	24
G06F9/4401	/	9	/	4401
G06F21/53	/	21	/	53
H04L9/08	/	9	/	08
H04L12/715	/	12	/	715
H04L12/771	/	12	/	771
G06F8/60	/	8	/	60
H04L29/06	/	29	/	06
G06F21/44	/	21	/	44
H04W4/00	/	4	/	00
G06F21/57	/	21	/	57
G06F9/455	/	9	/	455
G06F8/70	/	8	/	70

NON-CLAIMED				
H04L	/	29	/	08
H04L	/	12	/	24
G06F	/	9	/	4401
G06F	/	21	/	53
H04L	/	9	/	08
H04L	/	12	/	715
H04L	/	12	/	771
G06F	/	8	/	60
H04L	/	29	/	06
G06F	/	21	/	44
H04W	/	4	/	00
G06F	/	21	/	57
G06F	/	9	/	455
G06F	/	8	/	70

NONE (Assistant Examiner)	(Date)	**Total Claims Allowed:** 12	
/BRENDAN Y HIGA/ Primary Examiner, Art Unit 2441 (Primary Examiner)	27 April 2020 (Date)	O.G. Print Claim(s) 1	O.G. Print Figure 3

U.S. Patent and Trademark Office — Part of Paper No.: 20200426

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0340

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	BRENDAN HIGA	2441

US ORIGINAL CLASSIFICATION	
CLASS	**SUBCLASS**

CROSS REFERENCES(S)						
CLASS	**SUBCLASS (ONE SUBCLASS PER BLOCK)**					

NONE		**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	12	
/BRENDAN Y HIGA/ Primary Examiner, Art Unit 2441	27 April 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	3

U.S. Patent and Trademark Office — Part of Paper No.: 20200426

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0341

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	15/836,824	Taaghol et al.
	Examiner	**Art Unit**
	BRENDAN HIGA	2441

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☑ T.D. ☐ R.1.47

CLAIMS

Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
1	1	10	16												
2	2	11	17												
3	5	12	18												
4	7														
5	11														
6	12														
7	13														
8	14														
9	15														

NONE		**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	12	
/BRENDAN Y HIGA/ Primary Examiner, Art Unit 2441	27 April 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	3

U.S. Patent and Trademark Office — Part of Paper No.: 20200426

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0342

Bibliographic Data

Application No: 15/836,824

Foreign Priority claimed: ○ Yes ◉ No

35 USC 119 (a-d) conditions met: ☐ Yes ☑ No ☐ Met After Allowance

Verified and Acknowledged: /BRENDAN Y HIGA/ (Examiner's Signature) | BH (Initials)

Title: DISTRIBUTED SOFTWARE DEFINED NETWORKING

FILING or 371(c) DATE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
12/09/2017	709	2441	TAAG001CONT1
RULE			

APPLICANTS

Big Data Federation, Inc., Santa Clara, CA,

INVENTORS

Pouya Taaghol San Jose, CA, UNITED STATES

Vivek Ramanna San Jose, CA, UNITED STATES

CONTINUING DATA

This application is a CON of 14295331 06/04/2014 PAT 9843624

14295331 has PRO of 61834807 06/13/2013

FOREIGN APPLICATIONS

IF REQUIRED, FOREIGN LICENSE GRANTED**

01/10/2018

**** SMALL ENTITY ****

STATE OR COUNTRY

UNITED STATES

ADDRESS

Roark IP
1438 Dahlia Loop
San Jose, CA 95126
UNITED STATES

FILING FEE RECEIVED

$730

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L7	4877921	(container tenant mult$1tenant sandbox$4 virtual adj2 machine)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/27 09:50
L8	9084	(VM cloud server) adj25 (application web$1app app) same (open$1flow $3SDN software adj2 network$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/27 09:50
L9	72547	(VM cloud server) adj25 (application web$1app app cloud$1app) same L7	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/27 09:50
L10	27	(portal merchant market$1place) adj15 (application web$1app cloud$1app app) same L9 and L8	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/27 09:50
L11	20	L10 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2020/04/27 09:50

EAST Search History (Interference)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	5714	(VM cloud server) adj25 (application web$1app app) same (virtual$5 hyper$2visor open$1flow $3SDN software adj2 network$3).clm.	US-PGPUB; USPAT	OR	ON	2020/04/27 09:50

L2	2447	(portal web web$1sit$3 web$1pag$3 store merchant market$1place) adj25 (application web$1app app) adj10 (download$3 installation install$3).clm.	US-PGPUB; USPAT	OR	ON	2020/04/27 09:50
L3	84	L1 and L2	US-PGPUB; USPAT	OR	ON	2020/04/27 09:50
L4	70	L3 and (@pd @ptad @prad @ad @rlad)<"20130613"	US-PGPUB; USPAT	OR	ON	2020/04/27 09:50
L5	17393	(version$3 updat$3 recent$3 up$1grad$3) adj15 (authenticat$3 password authentic key check$1sum authoriz$6).clm.	US-PGPUB; USPAT	OR	ON	2020/04/27 09:50
L6	0	L4 and L5	US-PGPUB; USPAT	OR	ON	2020/04/27 09:50

4/27/2020 10:49:20 AM
C:\Users\bhiga\Documents\EAST\Workspaces\15836824_1.wsp

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), by mail or fax, or via EFS-Web.

By mail, send to: Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

By fax, send to: (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

112449 7590 05/04/2020
Roark IP
1438 Dahlia Loop
San Jose, CA 95126

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being transmitted to the USPTO via EFS-Web or by facsimile to (571) 273-2885, on the date below.

(Typed or printed name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

TITLE OF INVENTION: DISTRIBUTED SOFTWARE DEFINED NETWORKING

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$500	$0.00	$0.00	$500	08/04/2020

EXAMINER	ART UNIT	CLASS-SUBCLASS
HIGA, BRENDAN Y	2441	709-201000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❑ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

❑ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-09 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 ROARK IP
2
3

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document must have been previously recorded, or filed for recordation, as set forth in 37 CFR 3.11 and 37 CFR 3.81(a). Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE: BIG DATA FEDERATION, INC.

(B) RESIDENCE: (CITY and STATE OR COUNTRY): Santa Clara, CA

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ☒ Corporation or other private group entity ❑ Government

4a. Fees submitted: ☒ Issue Fee ❑ Publication Fee (if required) ❑ Advance Order - # of Copies _____

4b. Method of Payment: *(Please first reapply any previously paid fee shown above)*

❑ Electronic Payment via EFS-Web ❑ Enclosed check ❑ Non-electronic payment by credit card (Attach form PTO-2038)

❑ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment to Deposit Account No. _____

5. **Change in Entity Status** (from status indicated above)

❑ Applicant certifying micro entity status. See 37 CFR 1.29

❑ Applicant asserting small entity status. See 37 CFR 1.27

❑ Applicant changing to regular undiscounted fee status.

NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature /Christopher J. Horgan/ Date 5/6/2020

Typed or printed name Christopher J. Horgan Registration No. 40,394

PTOL-85 Part B (08-18) Approved for use through 01/31/2020 OMB 0651-0033 U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0346

Electronic Patent Application Fee Transmittal

Application Number:	15836824
Filing Date:	09-Dec-2017
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Filer:	Christopher James Horgan
Attorney Docket Number:	TAAG001CONT1

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
UTILITY APPL ISSUE FEE	2501	1	500	500

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**500**

Electronic Acknowledgement Receipt

EFS ID:	39369001
Application Number:	15836824
International Application Number:	
Confirmation Number:	9801
Title of Invention:	DISTRIBUTED SOFTWARE DEFINED NETWORKING
First Named Inventor/Applicant Name:	Pouya Taaghol
Customer Number:	112449
Filer:	Christopher James Horgan
Filer Authorized By:	
Attorney Docket Number:	TAAG001CONT1
Receipt Date:	06-MAY-2020
Filing Date:	09-DEC-2017
Time Stamp:	15:43:59
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	CARD
Payment was successfully received in RAM	$500
RAM confirmation Number	E202056F45384792
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Issue Fee Payment (PTO-85B)	TAAG001CONT1_ISSUE_FEE.pdf	90815 a4c808ab14227b103d4a9e1c9895edcbbd80748a	no	1
Warnings:					
Information:					
2	Fee Worksheet (SB06)	fee-info.pdf	30514 69ee1843f711a3971e17eb9df527c8a5d093b624	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	121329		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.
National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.
New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PRINTER RUSH
(PTO ASSISTANCE)

Application: 15836824	**Examiner:** Higa	**GAU:** 2441
From:	**Location:** IDC	**Creation Date:** 05/11/2020

Week Date: 5/4/2020

DOC CODE	DOC DATE	MISCELLANEOUS
[X] 1449	06/25/2019	[ ] Continuing Data
[ ] IDS		[ ] Foreign Priority
[ ] CLM		[ ] Document Legibility
[ ] IIFW/FWCLM		[ ] Fees
[ ] SRFW		[ ] Petition (TC)
[ ] DRW		[ ] Other
[ ] OATH		
[ ] 312		
[ ] SPEC		

[RUSH] Message:

1449's 6/25/2019 on page 1 of 3, cite no. 16 contains 2009027677 which is an incomplete Pre-Grant Application Publication document number. Please correct the document number and please confirm that you reviewed the document with that number.

Thank you.

[XRUSH] Response:

Initials: _____

Examiner: PUBS contact - James Leija, (703) 756-1197
NOTE: This form will be included as part of the official USPTO record with the response document coded as XRUSH.
REV: AUG 16

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0351

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	12/09/2017	Pouya Taaghol	TAAG001CONT1	9801

112449 7590 05/14/2020

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

EXAMINER	
HIGA, BRENDAN Y	

ART UNIT	PAPER NUMBER
2441	

MAIL DATE	DELIVERY MODE
05/14/2020	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Supplemental ***Notice of Allowability***	**Application No.** 15/836,824	**Applicant(s)** Taaghol et al.	
	Examiner BRENDAN HIGA	**Art Unit** 2441	**AIA (FITF) Status** Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☑ This communication is responsive to Applicant's 12/09/2017 IDS.
   ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

3. ☑ The allowed claim(s) is/are 1-2,5,7 and 11-18 . As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information , please see **http://www.uspto.gov/patents/init_events/pph/index.jsp** or send an inquiry to **PPHfeedback@uspto.gov.**

4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   Certified copies:

   a) ☐ All b) ☐ Some *c) ☐ None of the:

   1. ☐ Certified copies of the priority documents have been received.
   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
   3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____ .

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____ .

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☐ Notice of References Cited (PTO-892)
2. ☑ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date 12/09/2017.
3. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material _____.
4. ☐ Interview Summary (PTO-413), Paper No./Mail Date. _____.
5. ☐ Examiner's Amendment/Comment
6. ☐ Examiner's Statement of Reasons for Allowance
7. ☐ Other _____.

/BRENDAN Y HIGA/
Primary Examiner, Art Unit 2441

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0353

PRINTER RUSH
(PTO ASSISTANCE)

Application: 15836824	**Examiner:** Higa	**GAU:** 2441
From:	**Location:** IDC	**Creation Date:** 05/11/2020

Week Date: 5/4/2020

DOC CODE	DOC DATE	MISCELLANEOUS
[X] 1449	06/25/2019	[ ] Continuing Data
[ ] IDS		[ ] Foreign Priority
[ ] CLM		[ ] Document Legibility
[ ] IIFW/FWCLM		[ ] Fees
[ ] SRFW		[ ] Petition (TC)
[ ] DRW		[ ] Other
[ ] OATH		
[ ] 312		
[ ] SPEC		

[RUSH] Message:

1449's 6/25/2019 on page 1 of 3, cite no. 16 contains 2009027677 which is an incomplete Pre-Grant Application Publication document number. Please correct the document number and please confirm that you reviewed the document with that number.

Thank you.

[XRUSH] Response:

Corrected and considered 1449.

Initials: /B.Y.H/

Examiner: PUBS contact - James Leija, (703) 756-1197
NOTE: This form will be included as part of the official USPTO record with the response document coded as XRUSH.
REV: AUG 16

Exhibit 1002
Microsoft Corporation v. Edge Networking Systems, LLC
0354

PTO/SB/08a (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 1 of 3

Complete if Known	
Application Number	15/836824
Filing Date	
First Named Inventor	Taaghol
Art Unit	
Examiner Name	
Attorney Docket Number	TAAGCONT1

U. S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code[2] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
		US-8,528,108	09-2013	Ammer, Gerhard	
		US-8,756,595	06-2014	Arcilla, Andre	
		US-9,226,333	12-2015	Steiner, Moritz	
		US-20090300057	12-2009	Friedman, Nathaniel	
		US-20140241158	08-2014	Anthony, Jr., Bruce	
		US-20110075611	03-2011	Choi, Hyung-Nam	
		US-7,630,706	12-2009	Weigand, David	
		US-20130007254	01-2013	Fries, Robert	
		US-20090113327	04-2009	Pogrebinsky, Vladimir	
		US-20080209491	08-2008	Hasek, Charles	
		US-20140317293	10-2014	Shatzkamer, Kevin	
		US-20140254373	09-2014	Varma, Subir	
		US-20140259012	09-2014	Nandiali, Vishwamitra	
		US-20150173121	06-2015	Miklos, Gyorgy	
		US-20070078988	04-2007	Miloushev, Vladimir	
/B.Y.H/		US-2009027677 1	11-2009	Nickolov, Peter	
		US-20020016954	02-2002	Charities, Dietrich	
		US-7114021	09-2006	Seshadri	
		US-8132166	03-2012	DeHaan	

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document Country Code[3]-Number[4]-Kind Code[5] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]

Examiner Signature	/BRENDAN Y HIGA/	Date Considered	05/13/2020

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /B.Y.H/

PTO/SB/08a (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 2 of 3

Complete if Known	
Application Number	15/836824
Filing Date	
First Named Inventor	Taaghol
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001CONT1

U. S. PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Document Number Number-Kind Code[2] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages or Relevant Figures Appear
		US-20040107417	06-2004	Chis	
		US-20050246703	11-2005	Ahonen	
		US-20090007091	01-2009	Appiah	
		US-20040237081	11-2004	Homilier, Daniel	
		US-20120300615	11-2012	Kempf, James	
		US-20130160147	06-2013	Draluk, Vadim	
		US-20130263104	10-2013	Baset, Salman	
		US-20140062976	02-2014	Marino, Joseph	
		US-8,650,299	02-2014	Huang, Alex	
		US-8,843,933	09-2014	Holler, Anne Marie	
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			
		US-			

FOREIGN PATENT DOCUMENTS

Examiner Initials*	Cite No.[1]	Foreign Patent Document Country Code[3]-Number[4]-Kind Code[5] *(if known)*	Publication Date MM-DD-YYYY	Name of Patentee or Applicant of Cited Document	Pages, Columns, Lines, Where Relevant Passages Or Relevant Figures Appear	T[6]

Examiner Signature	/BRENDAN Y HIGA/	Date Considered	05/13/2020

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant. [1] Applicant's unique citation designation number (optional). [2] See Kinds Codes of USPTO Patent Documents at www.uspto.gov or MPEP 901.04. [3] Enter Office that issued the document, by the two-letter code (WIPO Standard ST.3). [4] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [5] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [6] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /B.Y.H/

PTO/SB/08b (07-09)
Approved for use through 11/30/2020. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Substitute for form 1449/PTO

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Use as many sheets as necessary)

Sheet 3 of 3

	Complete if Known
Application Number	15/836824
Filing Date	
First Named Inventor	TAAGHOL
Art Unit	
Examiner Name	
Attorney Docket Number	TAAG001CONT1

NON PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No.[1]	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume-issue number(s), publisher, city and/or country where published.	T[2]
		Bansal, et al., "OpenRadio: A Programmable Wireless Dataplane", published in HotSDN'12, Aug. 13, 2012, Helsinki, Finland	
		Stallings et al., "Software-Defined Networks and Open Flow" published in The Internet Protocol Journal, Vol. 16, No. 1, March 2012	
		Luo et al., "Software Defined Networking across Distributed Datacenters over Cloud" published 2013 IEEE International Conference on Cloud Computing Technology and Science	
		McKeown et al., "OpenFlow:Enabling Innovation in Campus Networks", March 14, 2008.	
		Li et al., "Toward Software-Defined Cellular Networks", March 14, 2008.	
		"Network Functions Visualization", White Paper published Oct. 22-24, 2012, at the "SDN and OpenFlow World Congress", Damstadt-Germany	
		Yap et al., "OpenRoads: Empowering Research in Mobile Networks" published in 2010 ACM SIGCOMM Computer Communication Review	
		"A Short Introduction To Cloud Platforms " by David Chappell, Aug 2008	

Examiner Signature	/BRENDAN Y HIGA/	Date Considered	05/13/2020

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] Applicant's unique citation designation number (optional). [2] Applicant is to place a check mark here if English language Translation is attached.

This collection of information is required by 37 CFR 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 2 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 (1-800-786-9199) and select option 2.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /B.Y.H/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/836,824	06/16/2020	10686871	TAAG001CONT1	9801

112449 7590 05/27/2020

Roark IP
1438 Dahlia Loop
San Jose, CA 95126

ISSUE NOTIFICATION

The projected patent number and issue date are specified above.

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment is 106 day(s). Any patent to issue from the above-identified application will include an indication of the adjustment on the front page.

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (http://pair.uspto.gov).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Application Assistance Unit (AAU) of the Office of Data Management (ODM) at (571)-272-4200.

APPLICANT(s) (Please see PAIR WEB site http://pair.uspto.gov for additional applicants):

Pouya Taaghol, San Jose, CA;
Big Data Federation, Inc., Santa Clara, CA;
Vivek Ramanna, San Jose, CA;

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The USA offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to encourage and facilitate business investment. To learn more about why the USA is the best country in the world to develop technology, manufacture products, and grow your business, visit SelectUSA.gov.