Characterising the IETF Through the Lens of RFC Deployment

Stephen McQuistin University of Glasgow
Mladen Karan Queen Mary University of London
Prashant Khare Queen Mary University of London
Colin Perkins University of Glasgow
Gareth Tyson Queen Mary University of London
Matthew Purver Queen Mary University of London
Patrick Healey Queen Mary University of London
Waleed Iqbal Queen Mary University of London
Junaid Qadir Information Technology University
Ignacio Castro Queen Mary University of London

ACM Internet Measurement Conference 2021

sodestream.github.io

Engineering and Physical Sciences Research Council

This work is funded by the UK Engineering and Physical Sciences Research Council, under grants EP/S033564/1 and EP/S036075/1.

Network Working Group
Request for Comments: 4535
Category: Standards Track

H. Harney
U. Meth
A. Colegrove
SPARTA, Inc.
G. Gross
IdentAware
June 2006

GSAKMP: Group Secure Association Key Management Protocol

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document specifies the Group Secure Association Key Management Protocol (GSAKMP). The GSAKMP provides a security framework for creating and managing cryptographic groups on a network. It provides mechanisms to disseminate group policy and authenticate users, rules to perform access control decisions during group establishment and recovery, capabilities to recover from the compromise of group members, delegation of group security functions, and capabilities to destroy the group. It also generates group keys.

Network Working Group
Request for Comments: 854

J. Postel
J. Reynolds
ISI

Obsoletes: NIC 18639

May 1983

TELNET PROTOCOL SPECIFICATION

This RFC specifies a standard for the ARPA Internet community. Hosts on the ARPA Internet are expected to adopt and implement this standard.

INTRODUCTION

The purpose of the TELNET Protocol is to provide a fairly general, bi-directional, eight-bit byte oriented communications facility. Its primary goal is to allow a standard method of interfacing terminal devices and terminal-oriented processes to each other. It is envisioned that the protocol may also be used for terminal-terminal communication ("linking") and process-process communication (distributed computation).

GENERAL CONSIDERATIONS

A TELNET connection is a Transmission Control Protocol (TCP) connection used to transmit data with interspersed TELNET control information.

The TELNET Protocol is built upon three main ideas: first, the concept of a "Network Virtual Terminal"; second, the principle of negotiated options; and third, a symmetric view of terminals and processes.

1. When a TELNET connection is first established, each end is assumed to originate and terminate at a "Network Virtual Terminal", or NVT. An NVT is an imaginary device which provides a standard, network-wide, intermediate representation of a canonical terminal. This eliminates the need for "server" and "user" hosts to keep information about the characteristics of each other's terminals and terminal handling conventions. All hosts, both user and server, map their local device characteristics and conventions so as to appear to be dealing with an NVT over the network, and each can assume a similar mapping by the other party. The NVT is intended to strike a balance between being overly restrictive (not providing hosts a high

mapping into their loca
e (penalizing users with

host is the host to whic
ed, and the "server" hos
some service. As an al

Network Working Group
Request for Comments: 5090
Obsoletes: 4590
Category: Standards Track

B. Sterman
Kayote Networks
D. Sadolevsky
SecureOL, Inc.
D. Schwartz
Kayote Networks
D. Williams
Cisco Systems
W. Beck
Deutsche Telekom AG
February 2008

RADIUS Extension for Digest Authentication

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document defines an extension to the Remote Authentication Dial-In User Service (RADIUS) protocol to enable support of Digest Authentication, for use with HTTP-style protocols like the Session Initiation Protocol (SIP) and HTTP.

Network Working Group
Request for Comments: 4340
Category: Standards Track

E. Kohler
UCLA
M. Handley
UCL
S. Floyd
ICIR
March 2006

Datagram Congestion Control Protocol (DCCP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

The Datagram Congestion Control Protocol (DCCP) is a transport protocol that provides bidirectional unicast connections of congestion-controlled unreliable datagrams. DCCP is suitable for applications that transfer fairly large amounts of data and that can benefit from control over the tradeoff between timeliness and reliability.

Table of Contents

1. Introduction 5
2. Design Rationale 6
3. Conventions and Terminology 7
on 8
.....9
.....9
.....9
e 10
.....10
.....10
.....11
.....12
Mechanisms 14

Network Working Group
Request for Comments: 1209

D. Piscitello
J. Lawrence
Bell Communications Research
March 1991

The Transmission of IP Datagrams over the SMDS Service

Status of this Memo

This memo defines a protocol for the transmission of IP and ARP packets over a Switched Multi-megabit Data Service Network configured as a logical IP subnetwork. This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This memo describes an initial use of IP and ARP in an SMDS service environment configured as a logical IP subnetwork, LIS (described below). The encapsulation method used is described, as well as various service-specific issues. This memo does not preclude subsequent treatment of the SMDS Service in configurations other than LIS; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents. This document considers only directly connected IP end-stations or routers; issues raised by MAC level bridging are beyond the scope of this paper.

Acknowledgment

This memo draws heavily in both concept and text from [4], written by Jon Postel and Joyce K. Reynolds of ISI and [5], written by David Katz of Merit, Inc. The authors would also like to acknowledge the contributions of the IP Over SMDS Service working group of the Internet Engineering Task Force.

Conventions

The following language conventions are used in the items of specification in this document:

- MUST, SHALL, or MANDATORY -- the item is an absolute requirement of the specification.

Internet Engineering Task Force (IETF)
Request for Comments: 7011
STD: 77
Obsoletes: 5101
Category: Standards Track
ISSN: 2070-1721

B. Claise, Ed.
Cisco Systems, Inc.
B. Trammell, Ed.
ETH Zurich
P. Aitken
Cisco Systems, Inc.
September 2013

Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information

Abstract

This document specifies the IP Flow Information Export (IPFIX) protocol, which serves as a means for transmitting Traffic Flow information over the network. In order to transmit Traffic Flow information from an Exporting Process to a Collecting Process, a common representation of flow data and a standard means of communicating them are required. This document describes how the IPFIX Data and Template Records are carried over a number of transport protocols from an IPFIX Exporting Process to an IPFIX Collecting Process. This document obsoletes RFC 5101.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7011.

Network Working Group
Request for Comments: 2741
Obsoletes: 2257
Category: Standards Track

M. Daniele
Compaq Computer Corporation
B. Wijnen
T.J. Watson Research Center, IBM Corp.
M. Ellison, Ed.
Ellison Software Consulting, Inc.
D. Francisco. Ed.
Cisco Systems, Inc.
January 2000

Agent Extensibility (AgentX) Protocol
Version 1

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

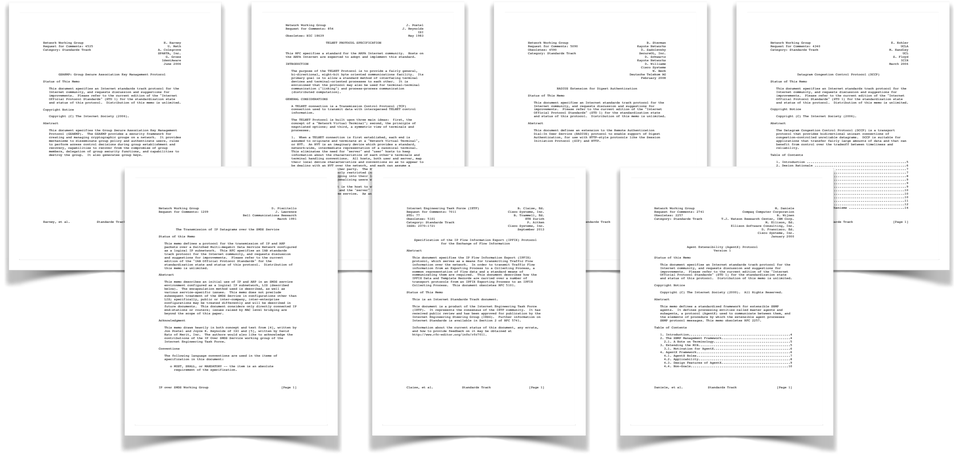
Abstract

This memo defines a standardized framework for extensible SNMP agents. It defines processing entities called master agents and subagents, a protocol (AgentX) used to communicate between them, and the elements of procedure by which the extensible agent processes SNMP protocol messages. This memo obsoletes RFC 2257.

Table of Contents

1. Introduction.....4
2. The SNMP Management Framework.....4
2.1. A Note on Terminology.....5
3. Extending the MIB.....5
3.1. Motivation for AgentX.....6
4. AgentX Framework.....6
4.1. AgentX Roles.....7
4.2. Applicability.....8
4.3. Design Features of AgentX.....9
4.4. Non-Goals.....10

sodestream.github.io



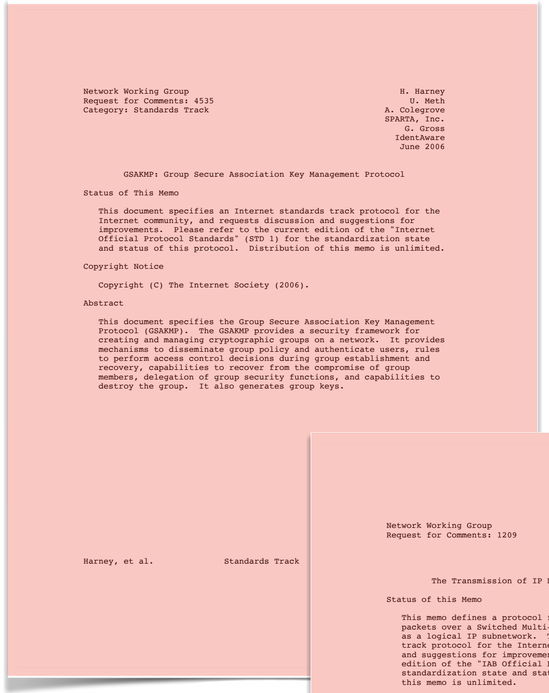
Documents





IETF®

sodestream.github.io



Network Working Group H. Harney
Request for Comments: 4535 U. Meth
Category: Standards Track A. Colegrove
SPARTA, Inc.
G. Gross
IdentAware
June 2006

GSAKMP: Group Secure Association Key Management Protocol

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

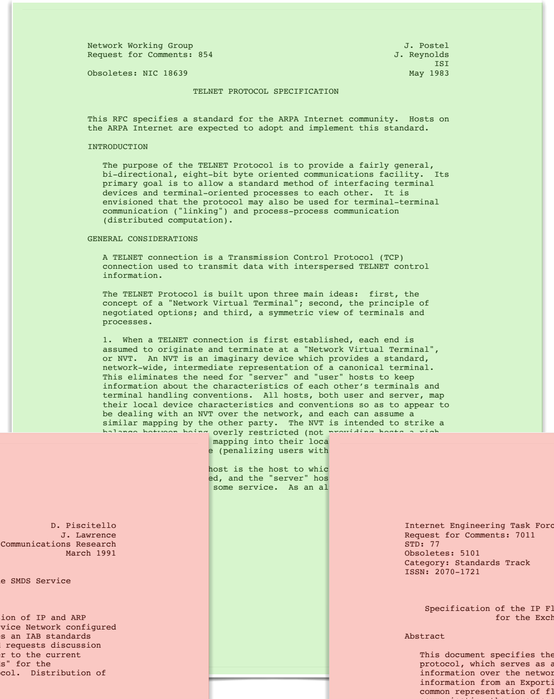
Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document specifies the Group Secure Association Key Management Protocol (GSAKMP). The GSAKMP provides a security framework for creating and managing cryptographic groups on a network. It provides mechanisms to disseminate group policy and authenticate users, rules to perform access control decisions during group establishment and recovery, capabilities to recover from the compromise of group members, delegation of group security functions, and capabilities to destroy the group. It also generates group keys.

Harney, et al. Standards Track



Network Working Group J. Postel
Request for Comments: 854 J. Reynolds
ISI
Obsoletes: NIC 18639 May 1983

TELNET PROTOCOL SPECIFICATION

This RFC specifies a standard for the ARPA Internet community. Hosts on the ARPA Internet are expected to adopt and implement this standard.

INTRODUCTION

The purpose of the TELNET Protocol is to provide a fairly general, bi-directional, eight-bit byte oriented communications facility. Its primary goal is to allow a standard method of interfacing terminal devices and terminal-oriented processes to each other. It is envisioned that the protocol may also be used for terminal-terminal communication ("linking") and process-process communication (distributed computation).

GENERAL CONSIDERATIONS

A TELNET connection is a Transmission Control Protocol (TCP) connection used to transmit data with interspersed TELNET control information.

The TELNET Protocol is built upon three main ideas: first, the concept of a "Network Virtual Terminal"; second, the principle of negotiated options; and third, a symmetric view of terminals and processes.

1. When a TELNET connection is first established, each end is assumed to originate and terminate at a "Network Virtual Terminal", or NVT. An NVT is an imaginary device which provides a standard, network-wide, intermediate representation of a canonical terminal. This eliminates the need for "server" and "user" hosts to keep information about the characteristics of each other's terminals and terminal handling conventions. All hosts, both user and server, map their local device characteristics and conventions so as to appear to be dealing with an NVT over the network, and each can assume a similar mapping by the other party. The NVT is intended to strike a



Network Working Group B. Sterman
Request for Comments: 5090 Kayote Networks
Obsoletes: 4590 D. Sadolevsky
Category: Standards Track SecureOL, Inc.
D. Schwartz
Kayote Networks
D. Williams
Cisco Systems
W. Beck
Deutsche Telekom AG
February 2008

RADIUS Extension for Digest Authentication

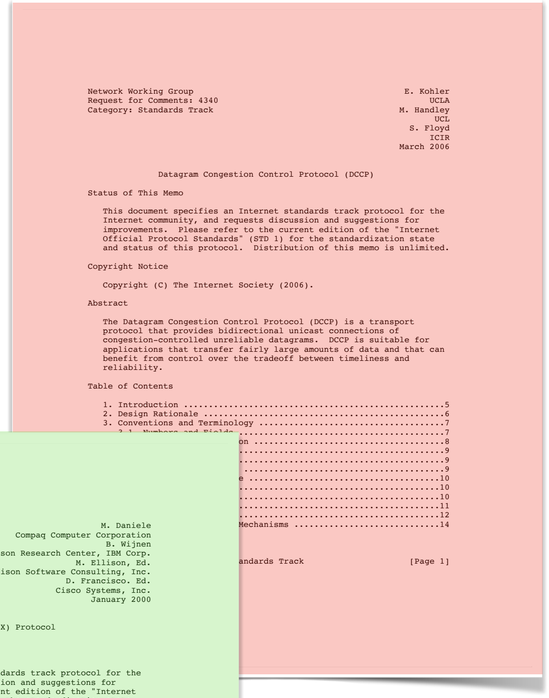
Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document defines an extension to the Remote Authentication Dial-In User Service (RADIUS) protocol to enable support of Digest Authentication, for use with HTTP-style protocols like the Session Initiation Protocol (SIP) and HTTP.

Standards Track



Network Working Group E. Kohler
Request for Comments: 4340 UCLA
Category: Standards Track M. Handley
UCL
S. Floyd
ICIR
March 2006

Datagram Congestion Control Protocol (DCCP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

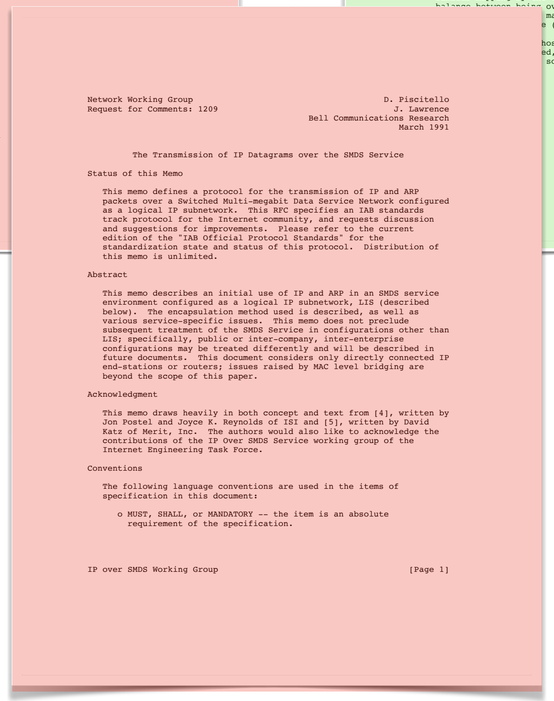
Abstract

The Datagram Congestion Control Protocol (DCCP) is a transport protocol that provides bidirectional unicast connections of congestion-controlled unreliable datagrams. DCCP is suitable for applications that transfer fairly large amounts of data and that can benefit from control over the tradeoff between timeliness and reliability.

Table of Contents

1. Introduction 5
2. Design Rationale 6
3. Conventions and Terminology 7
.....7
.....8
.....9
.....9
.....9
.....10
.....10
.....10
.....11
.....12
Mechanisms 14

andards Track [Page 1]



Network Working Group D. Piscitello
Request for Comments: 1209 J. Lawrence
Bell Communications Research
March 1991

The Transmission of IP Datagrams over the SMDS Service

Status of this Memo

This memo defines a protocol for the transmission of IP and ARP packets over a Switched Multi-megabit Data Service Network configured as a logical IP subnetwork. This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This memo describes an initial use of IP and ARP in an SMDS service environment configured as a logical IP subnetwork, LIS (described below). The encapsulation method used is described, as well as various service-specific issues. This memo does not preclude subsequent treatment of the SMDS Service in configurations other than LIS; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents. This document considers only directly connected IP end-stations or routers; issues raised by MAC level bridging are beyond the scope of this paper.

Acknowledgment

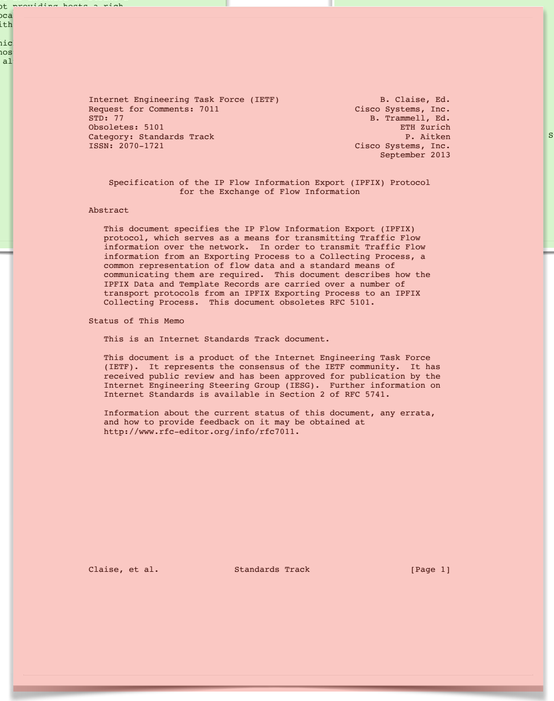
This memo draws heavily in both concept and text from [4], written by Jon Postel and Joyce K. Reynolds of ISI and [5], written by David Katz of Merit, Inc. The authors would also like to acknowledge the contributions of the IP Over SMDS Service working group of the Internet Engineering Task Force.

Conventions

The following language conventions are used in the items of specification in this document:

o MUST, SHALL, or MANDATORY -- the item is an absolute requirement of the specification.

IP over SMDS Working Group [Page 1]



Internet Engineering Task Force (IETF) B. Claise, Ed.
Request for Comments: 7011 Cisco Systems, Inc.
STD: 77 B. Trammell, Ed.
Obsoletes: 5101 ETH Zurich
Category: Standards Track P. Aitken
ISSN: 2070-1721 Cisco Systems, Inc.
September 2013

Specification of the IP Flow Information Export (IPFIX) Protocol
for the Exchange of Flow Information

Abstract

This document specifies the IP Flow Information Export (IPFIX) protocol, which serves as a means for transmitting Traffic Flow information over the network. In order to transmit Traffic Flow information from an Exporting Process to a Collecting Process, a common representation of flow data and a standard means of communicating them are required. This document describes how the IPFIX Data and Template Records are carried over a number of transport protocols from an IPFIX Exporting Process to an IPFIX Collecting Process. This document obsoletes RFC 5101.

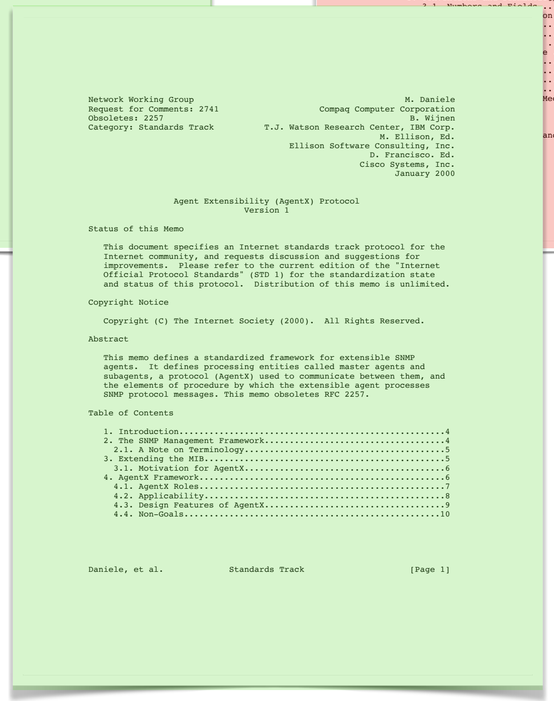
Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7011.

Claise, et al. Standards Track [Page 1]



Network Working Group M. Daniele
Request for Comments: 2741 Compaq Computer Corporation
Obsoletes: 2257 B. Wijnen
Category: Standards Track T.J. Watson Research Center, IBM Corp.
M. Ellison, Ed.
Ellison Software Consulting, Inc.
D. Francisco. Ed.
Cisco Systems, Inc.
January 2000

Agent Extensibility (AgentX) Protocol
Version 1

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This memo defines a standardized framework for extensible SNMP agents. It defines processing entities called master agents and subagents, a protocol (AgentX) used to communicate between them, and the elements of procedure by which the extensible agent processes SNMP protocol messages. This memo obsoletes RFC 2257.

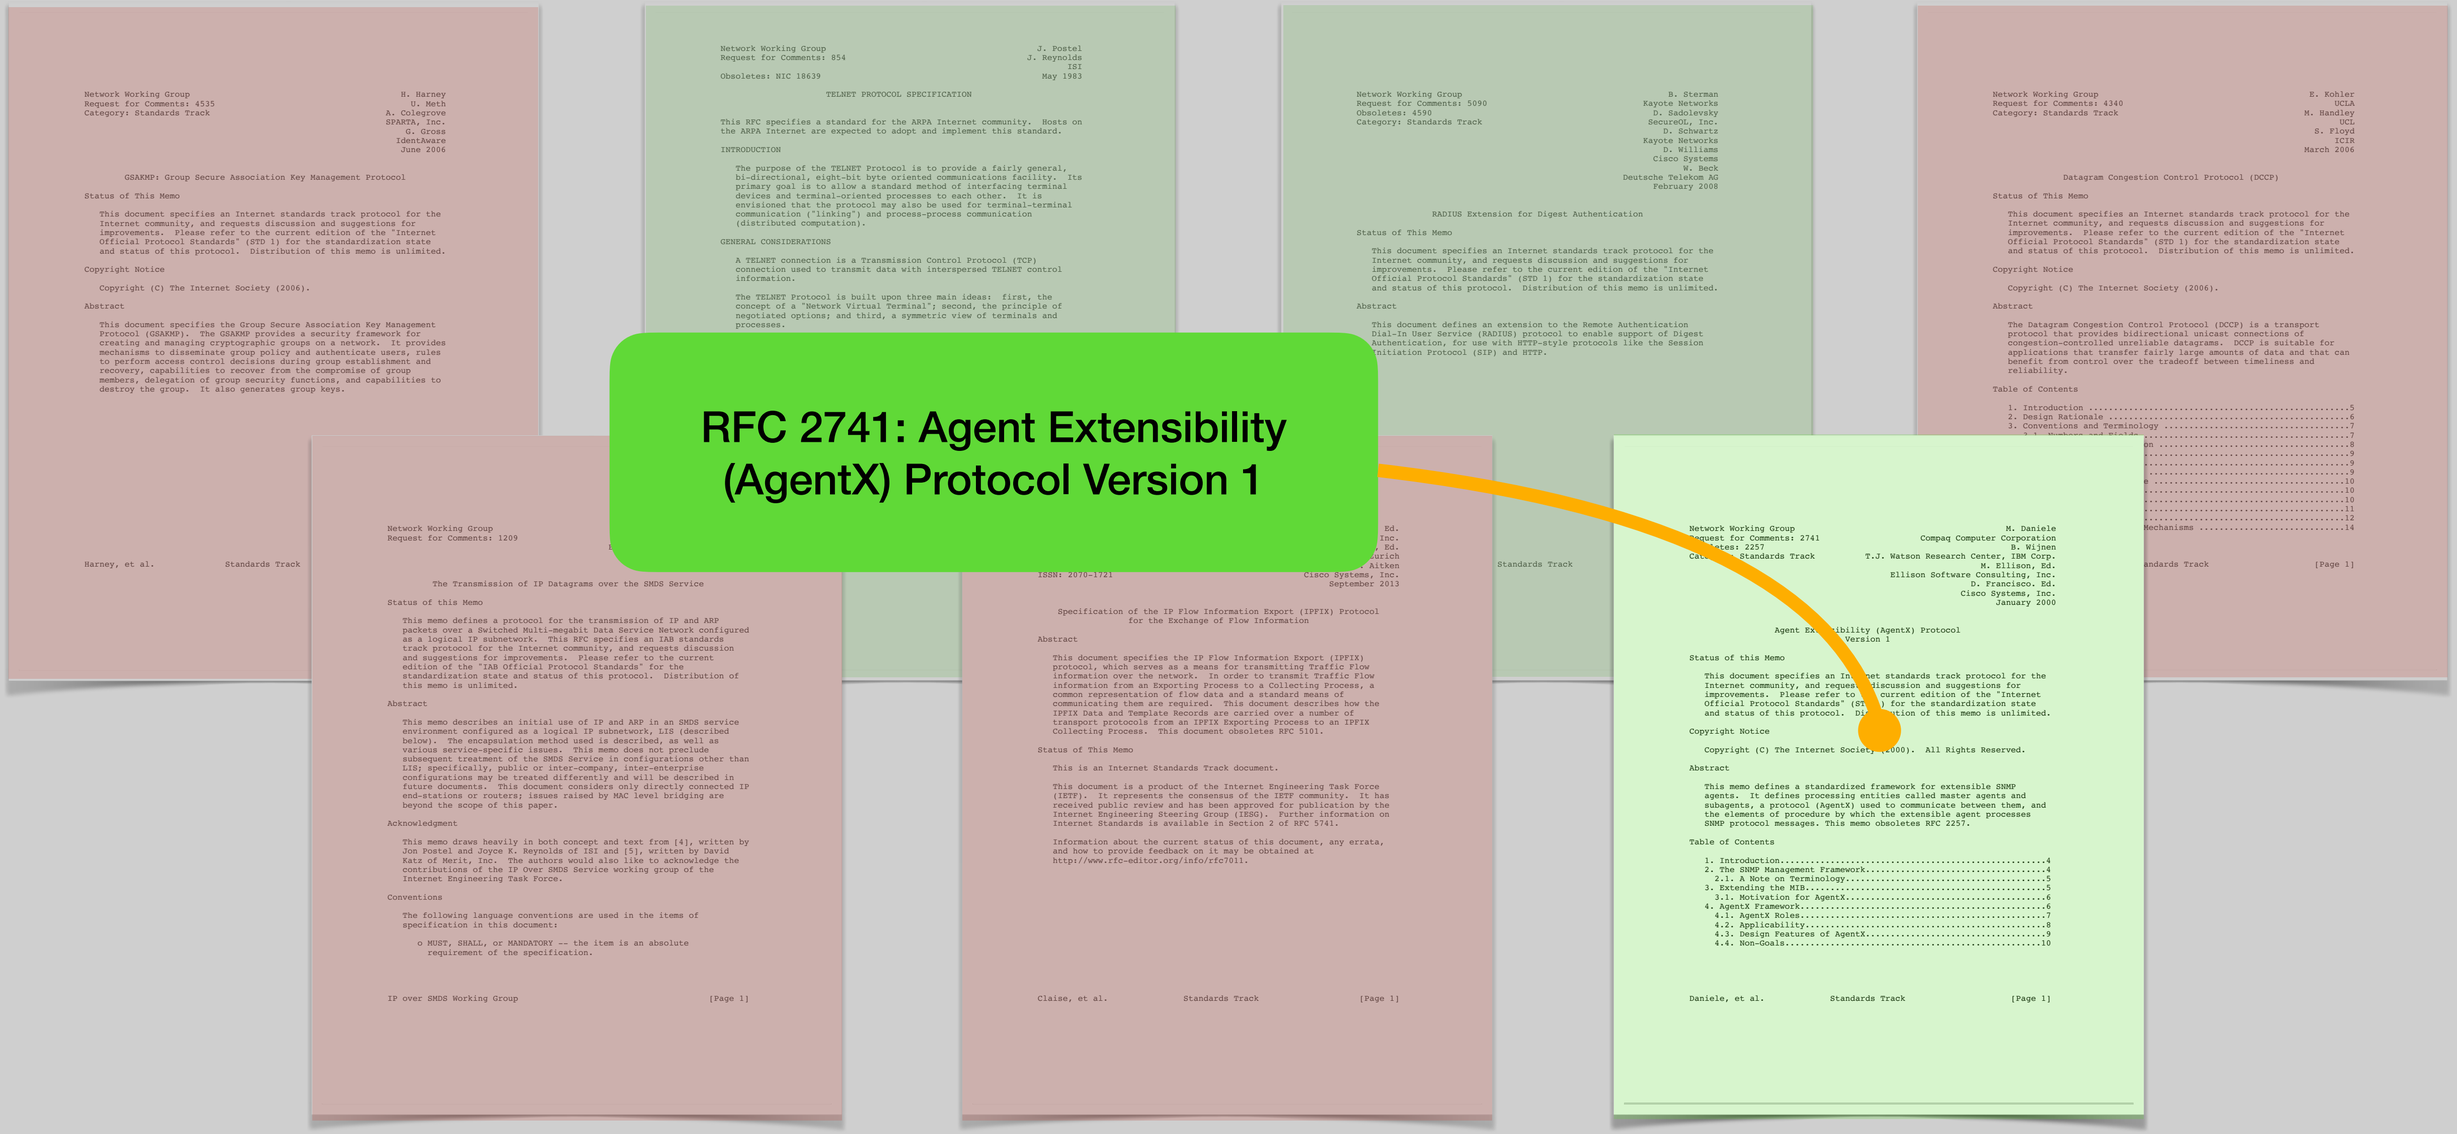
Table of Contents

1. Introduction.....4
2. The SNMP Management Framework.....4
2.1. A Note on Terminology.....5
3. Extending the MIB.....5
3.1. Motivation for AgentX.....6
4. AgentX Framework.....6
4.1. AgentX Roles.....7
4.2. Applicability.....8
4.3. Design Features of AgentX.....9
4.4. Non-Goals.....10

Daniele, et al. Standards Track [Page 1]

sodestream.github.io

RFC deployment data from Nikkhah et al., "*A Statistical Exploration of Protocol Adoption*". IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642



Network Working Group H. Harney
Request for Comments: 4535 U. Meth
Category: Standards Track A. Colegrove
SPARTA, Inc.
G. Gross
IdentAware
June 2006

GSAKMP: Group Secure Association Key Management Protocol

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document specifies the Group Secure Association Key Management Protocol (GSAKMP). The GSAKMP provides a security framework for creating and managing cryptographic groups on a network. It provides mechanisms to disseminate group policy and authenticate users, rules to perform access control decisions during group establishment and recovery, capabilities to recover from the compromise of group members, delegation of group security functions, and capabilities to destroy the group. It also generates group keys.

Harney, et al. Standards Track

Network Working Group J. Postel
Request for Comments: 854 J. Reynolds
ISI
Obsoletes: NIC 18639 May 1983

TELNET PROTOCOL SPECIFICATION

This RFC specifies a standard for the ARPA Internet community. Hosts on the ARPA Internet are expected to adopt and implement this standard.

INTRODUCTION

The purpose of the TELNET Protocol is to provide a fairly general, bi-directional, eight-bit byte oriented communications facility. Its primary goal is to allow a standard method of interfacing terminal devices and terminal-oriented processes to each other. It is envisioned that the protocol may also be used for terminal-terminal communication ("linking") and process-process communication (distributed computation).

GENERAL CONSIDERATIONS

A TELNET connection is a Transmission Control Protocol (TCP) connection used to transmit data with interspersed TELNET control information.

The TELNET Protocol is built upon three main ideas: first, the concept of a "Network Virtual Terminal"; second, the principle of negotiated options; and third, a symmetric view of terminals and processes.

Network Working Group B. Sterman
Request for Comments: 5090 Kayote Networks
Obsoletes: 4590 D. Sadolevsky
Category: Standards Track SecureOL, Inc.
D. Schwartz
Kayote Networks
D. Williams
Cisco Systems
W. Beck
Deutsche Telekom AG
February 2008

RADIUS Extension for Digest Authentication

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document defines an extension to the Remote Authentication Dial-In User Service (RADIUS) protocol to enable support of Digest Authentication, for use with HTTP-style protocols like the Session Initiation Protocol (SIP) and HTTP.

Network Working Group E. Kohler
Request for Comments: 4340 UCLA
Category: Standards Track M. Handley
UCL
S. Floyd
ICIR
March 2006

Datagram Congestion Control Protocol (DCCP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

The Datagram Congestion Control Protocol (DCCP) is a transport protocol that provides bidirectional unicast connections of congestion-controlled unreliable datagrams. DCCP is suitable for applications that transfer fairly large amounts of data and that can benefit from control over the tradeoff between timeliness and reliability.

Table of Contents

1. Introduction 5
2. Design Rationale 6
3. Conventions and Terminology 7

Standards Track [Page 1]

Network Working Group
Request for Comments: 1209

The Transmission of IP Datagrams over the SMDS Service

Status of This Memo

This memo defines a protocol for the transmission of IP and ARP packets over a Switched Multi-megabit Data Service Network configured as a logical IP subnetwork. This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This memo describes an initial use of IP and ARP in an SMDS service environment configured as a logical IP subnetwork, LIS (described below). The encapsulation method used is described, as well as various service-specific issues. This memo does not preclude subsequent treatment of the SMDS Service in configurations other than LIS; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents. This document considers only directly connected IP end-stations or routers; issues raised by MAC level bridging are beyond the scope of this paper.

Acknowledgment

This memo draws heavily in both concept and text from [4], written by Jon Postel and Joyce K. Reynolds of ISI and [5], written by David Katz of Merit, Inc. The authors would also like to acknowledge the contributions of the IP Over SMDS Service working group of the Internet Engineering Task Force.

Conventions

The following language conventions are used in the items of specification in this document:

o MUST, SHALL, or MANDATORY -- the item is an absolute requirement of the specification.

IP over SMDS Working Group [Page 1]

ISSN: 2070-1721 Cisco Systems, Inc.
September 2013

Specification of the IP Flow Information Export (IPFIX) Protocol
for the Exchange of Flow Information

Abstract

This document specifies the IP Flow Information Export (IPFIX) protocol, which serves as a means for transmitting Traffic Flow information over the network. In order to transmit Traffic Flow information from an Exporting Process to a Collecting Process, a common representation of flow data and a standard means of communicating them are required. This document describes how the IPFIX Data and Template Records are carried over a number of transport protocols from an IPFIX Exporting Process to an IPFIX Collecting Process. This document obsoletes RFC 5101.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7011.

Claise, et al. Standards Track [Page 1]

Network Working Group M. Daniele
Request for Comments: 2741 Compaq Computer Corporation
Obsoletes: 2257 B. Wijnen
Category: Standards Track T.J. Watson Research Center, IBM Corp.
M. Ellison, Ed.
Ellison Software Consulting, Inc.
D. Francisco. Ed.
Cisco Systems, Inc.
January 2000

Agent Extensibility (AgentX) Protocol
Version 1

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This memo defines a standardized framework for extensible SNMP agents. It defines processing entities called master agents and subagents, a protocol (AgentX) used to communicate between them, and the elements of procedure by which the extensible agent processes SNMP protocol messages. This memo obsoletes RFC 2257.

Table of Contents

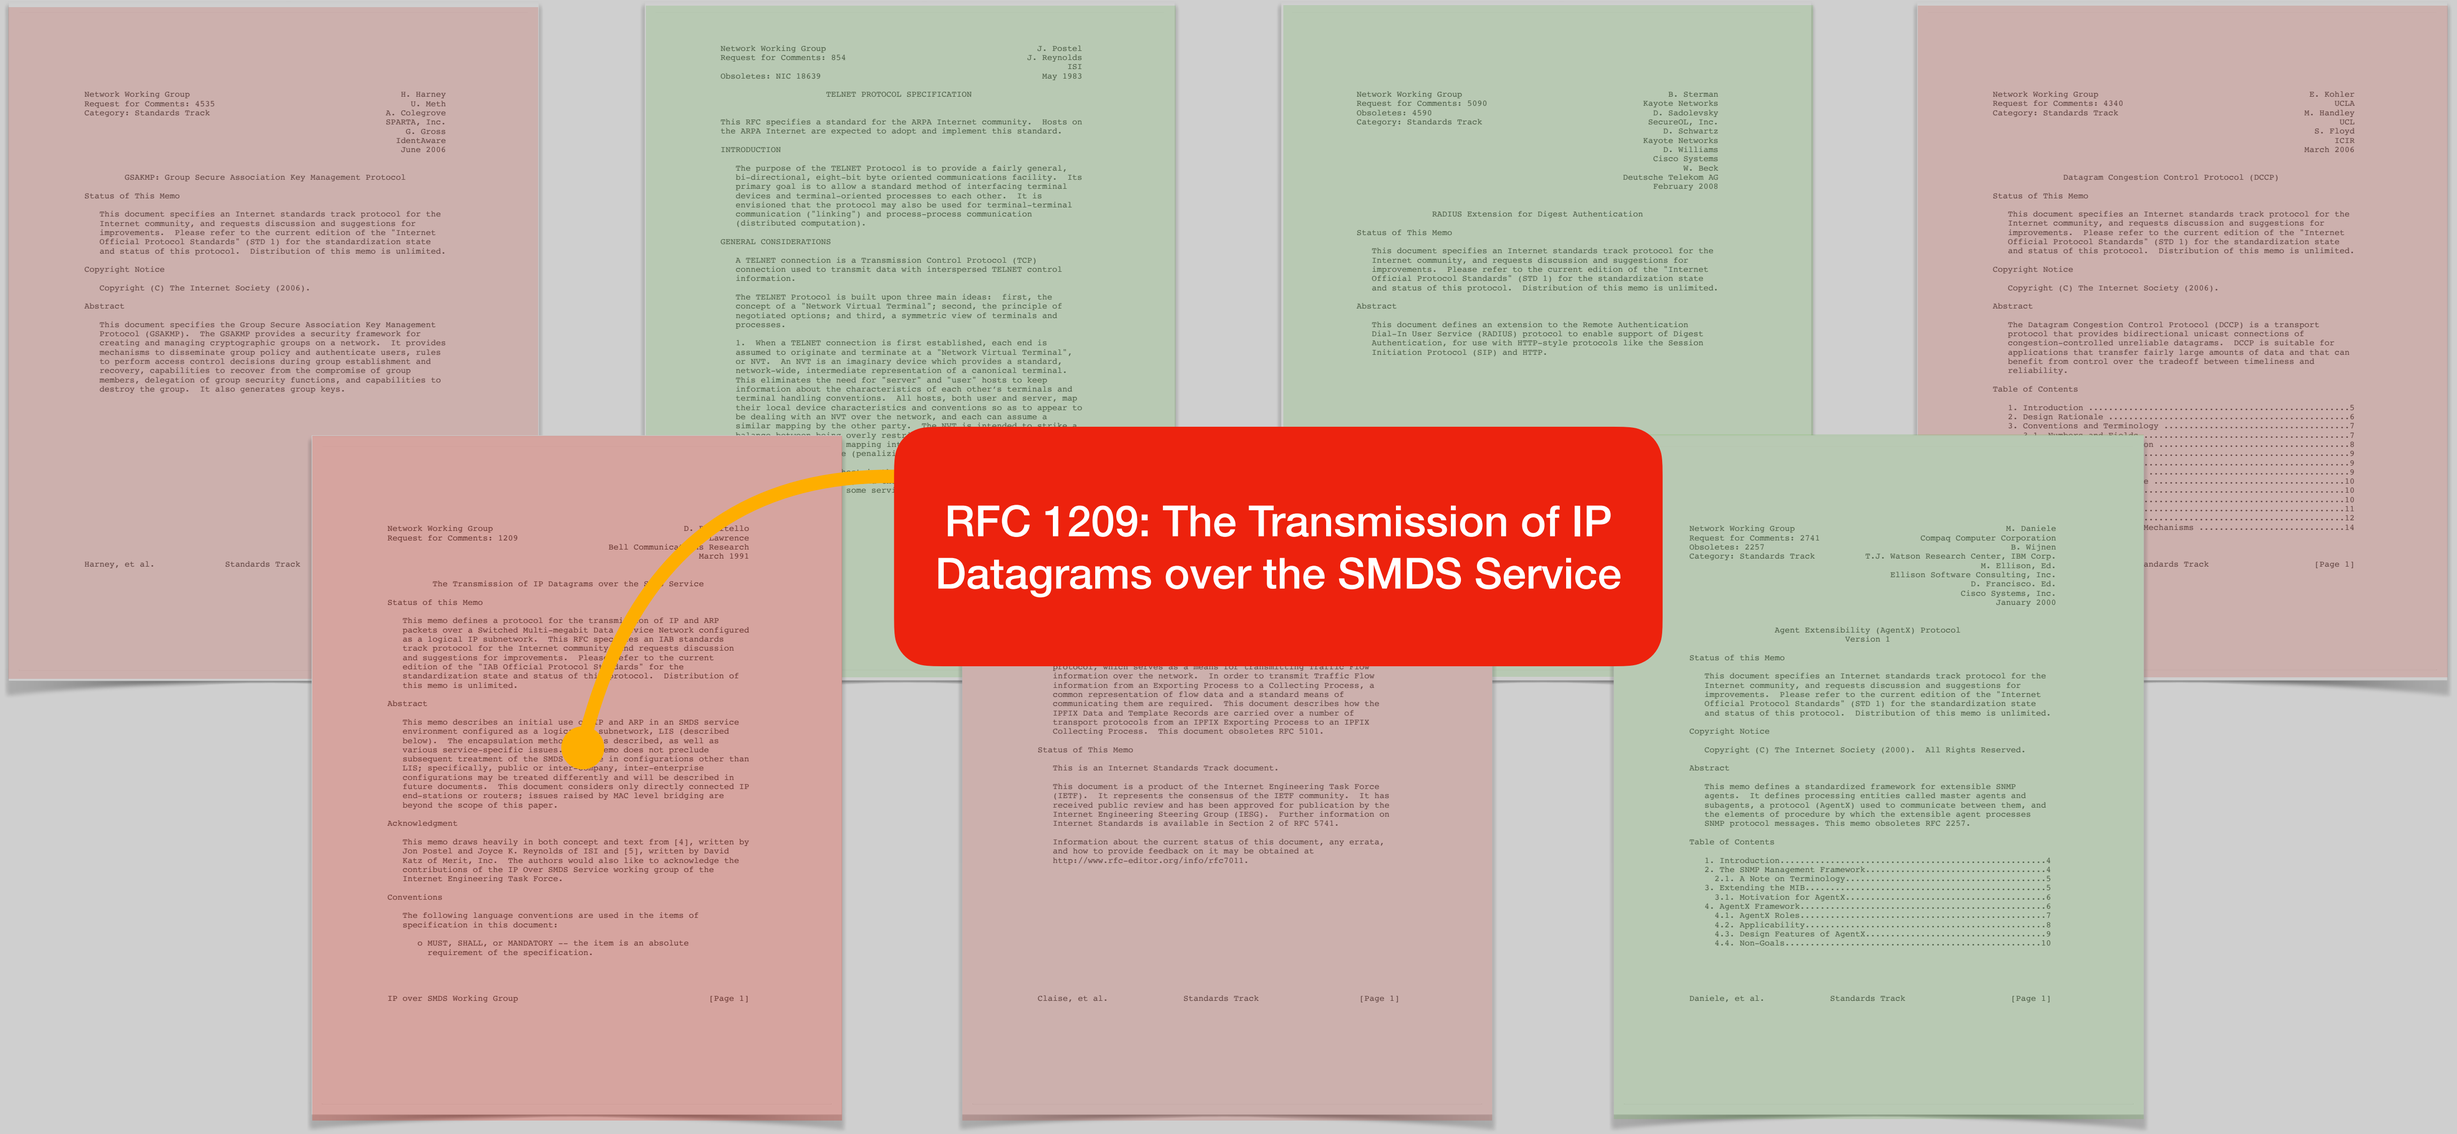
1. Introduction.....4
2. The SNMP Management Framework.....4
2.1. A Note on Terminology.....5
3. Extending the MIB.....5
3.1. Motivation for AgentX.....6
4. AgentX Framework.....6
4.1. AgentX Roles.....7
4.2. Applicability.....8
4.3. Design Features of AgentX.....9
4.4. Non-Goals.....10

Daniele, et al. Standards Track [Page 1]

RFC 2741: Agent Extensibility (AgentX) Protocol Version 1

sodestream.github.io

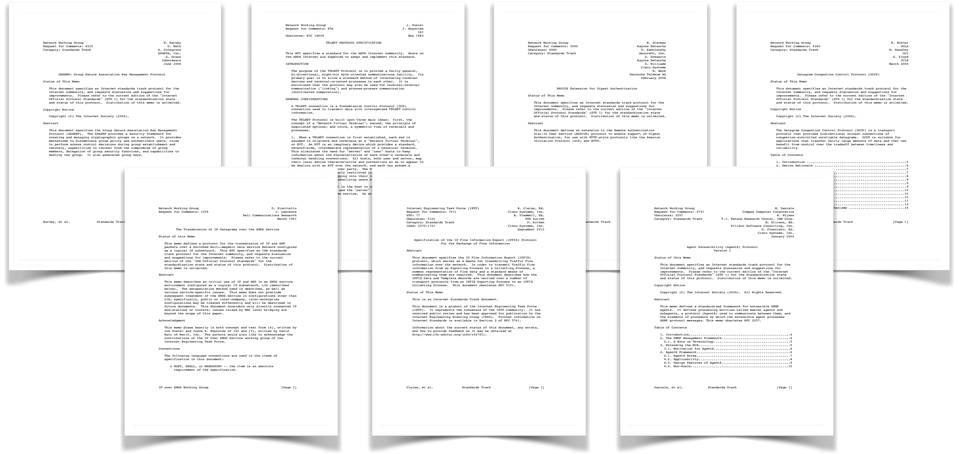
RFC deployment data from Nikkhah et al., "*A Statistical Exploration of Protocol Adoption*". IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642



RFC 1209: The Transmission of IP Datagrams over the SMDS Service

sodestream.github.io

RFC deployment data from Nikkhah et al., “*A Statistical Exploration of Protocol Adoption*”. IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642



Documents

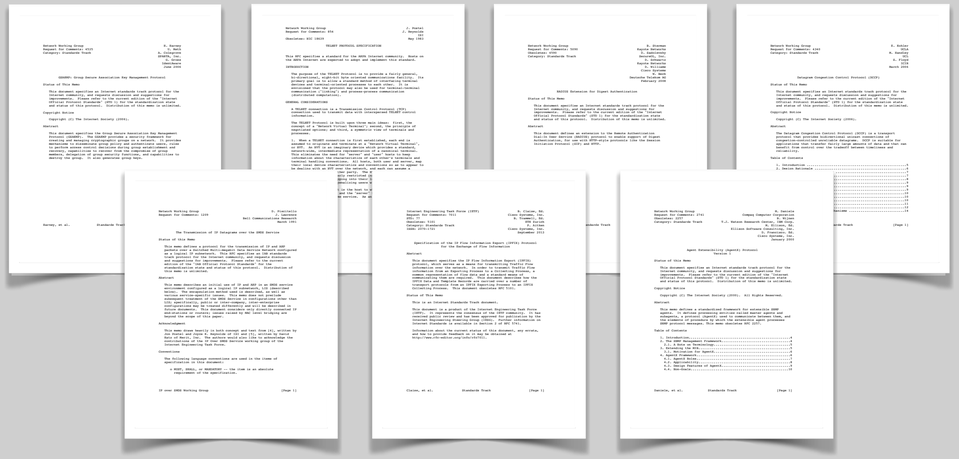
Authors

I E T F®

Mailing lists

sodestream.github.io

Nikkhah et al., “*A Statistical Exploration of Protocol Adoption*”.
IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642



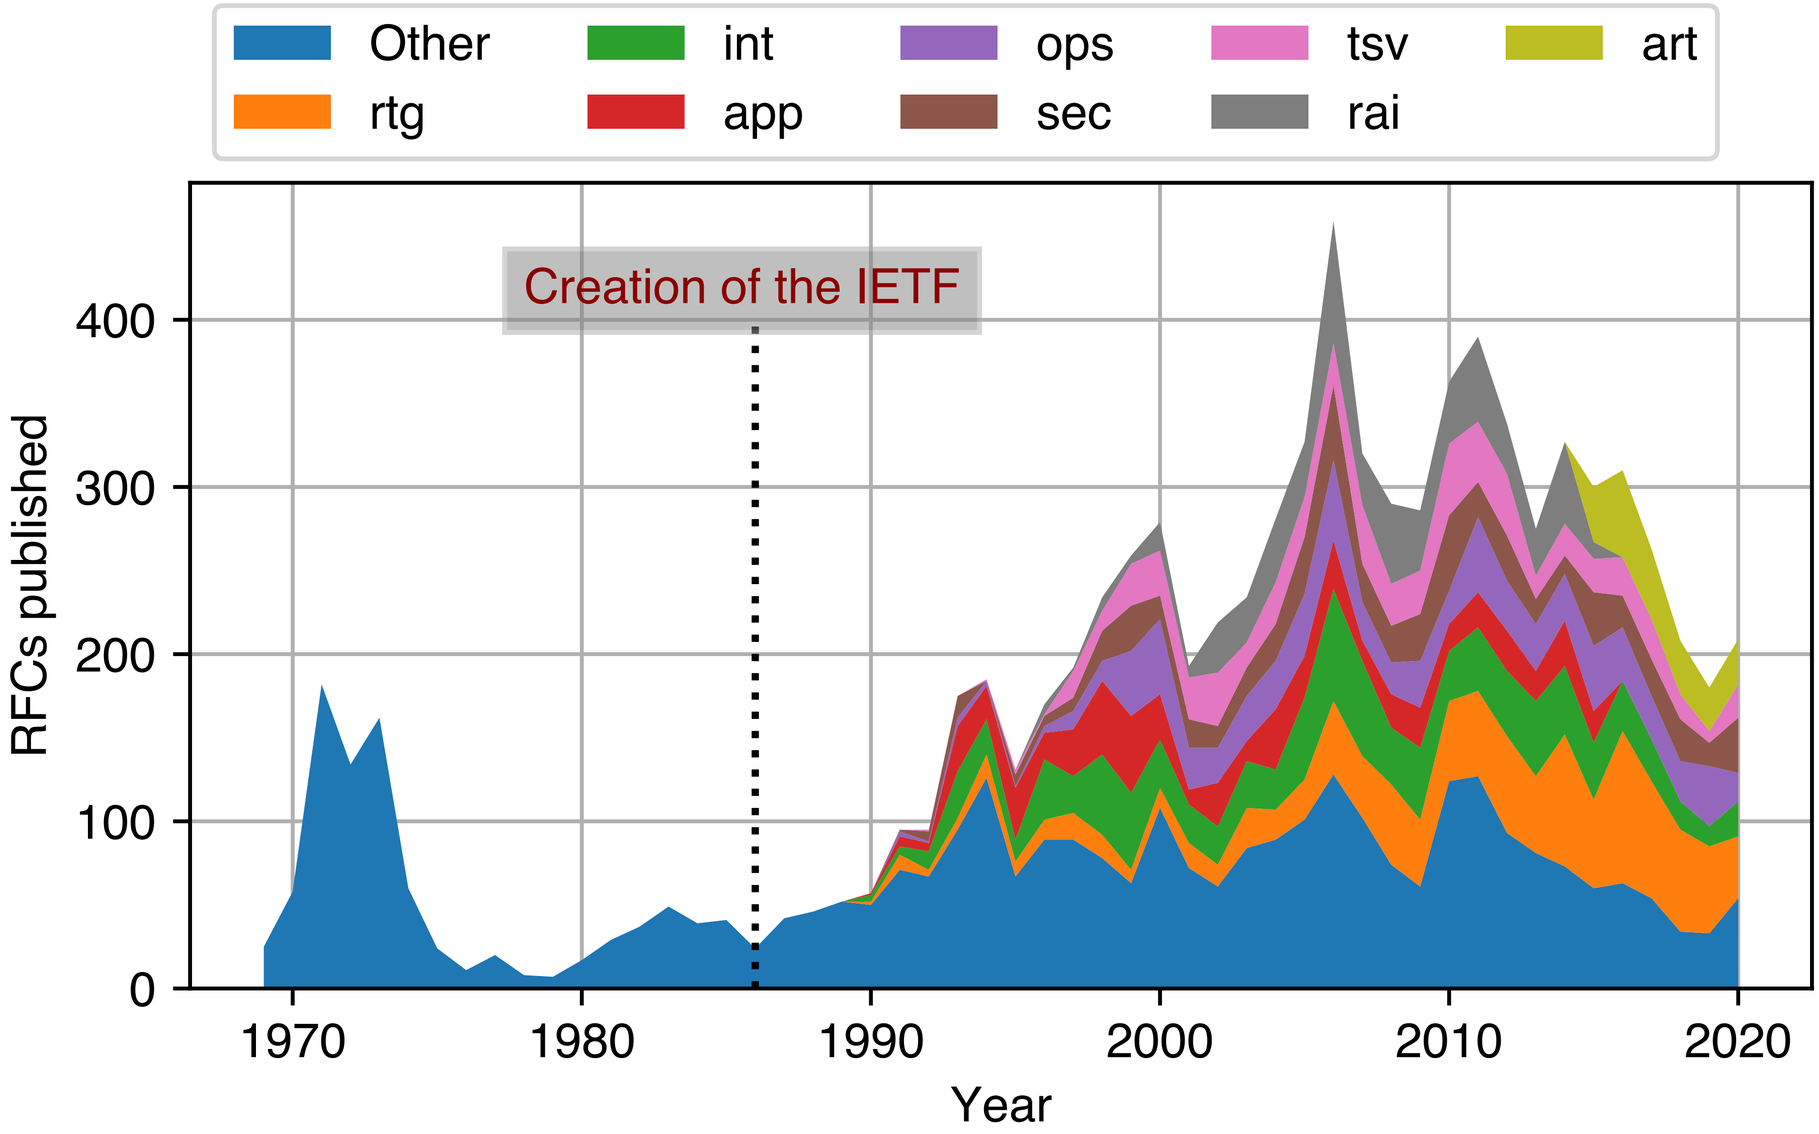
Documents

Authors

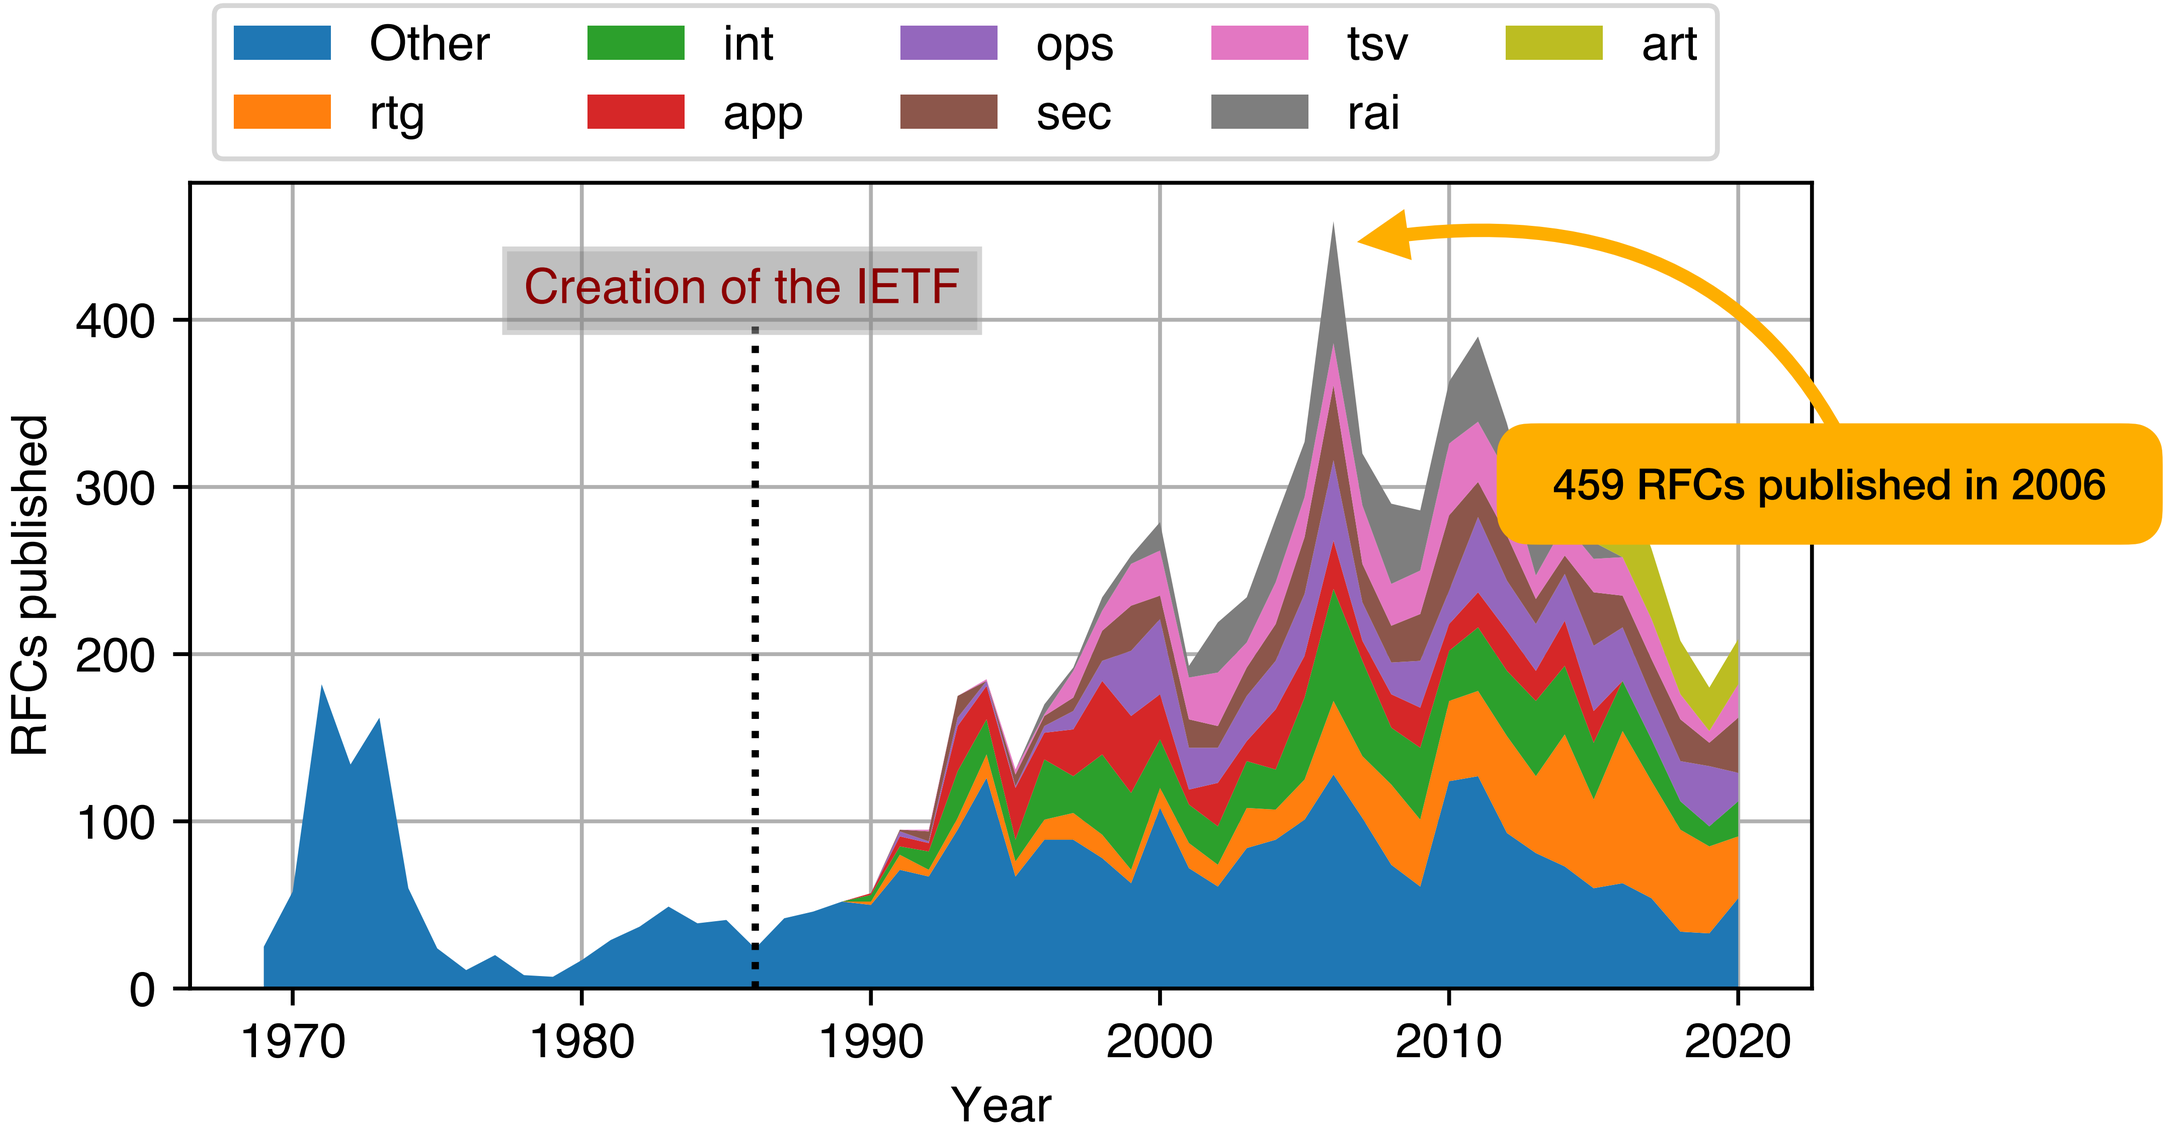
IETF®

Mailing lists

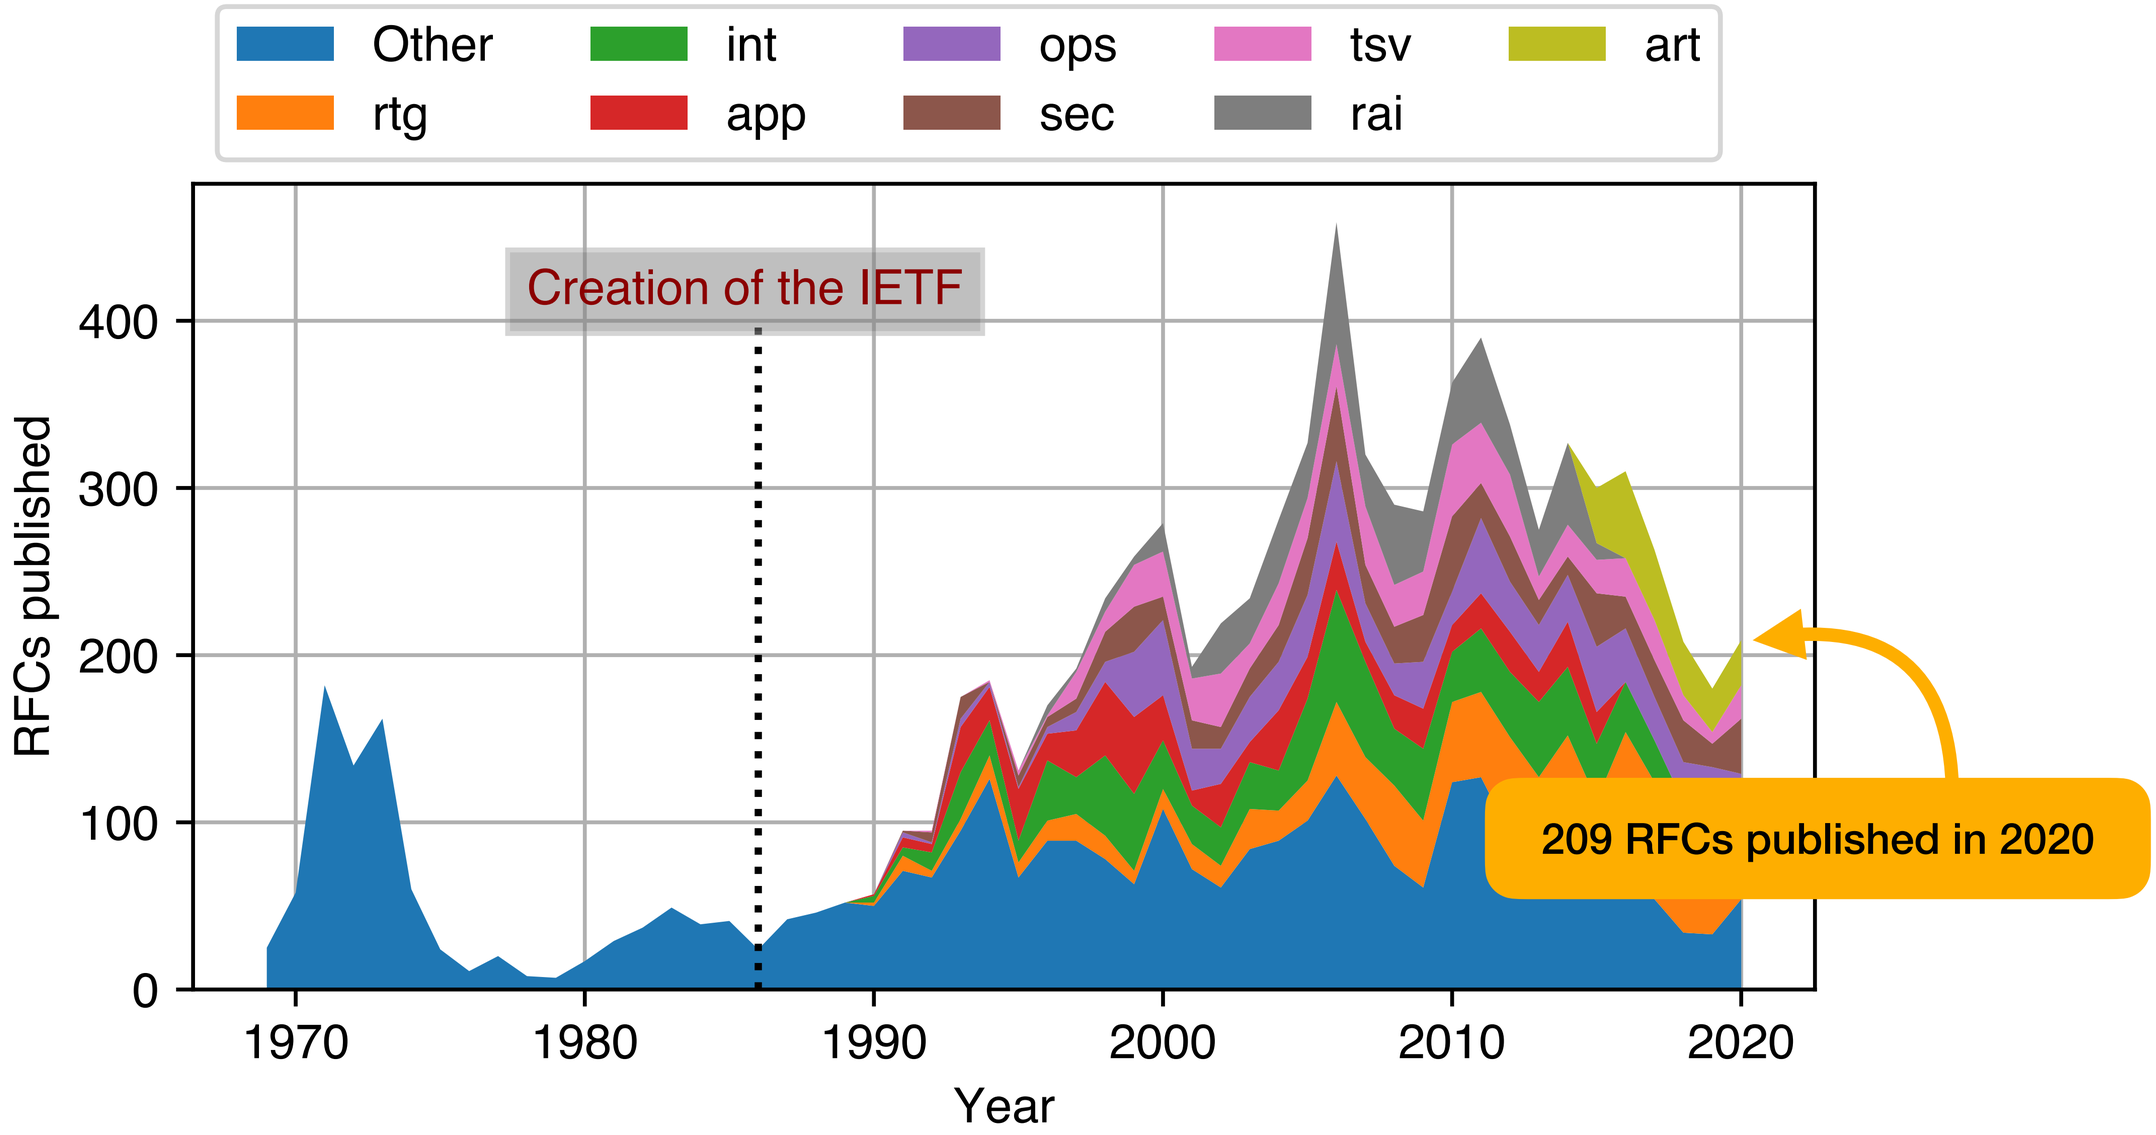
sodestream.github.io



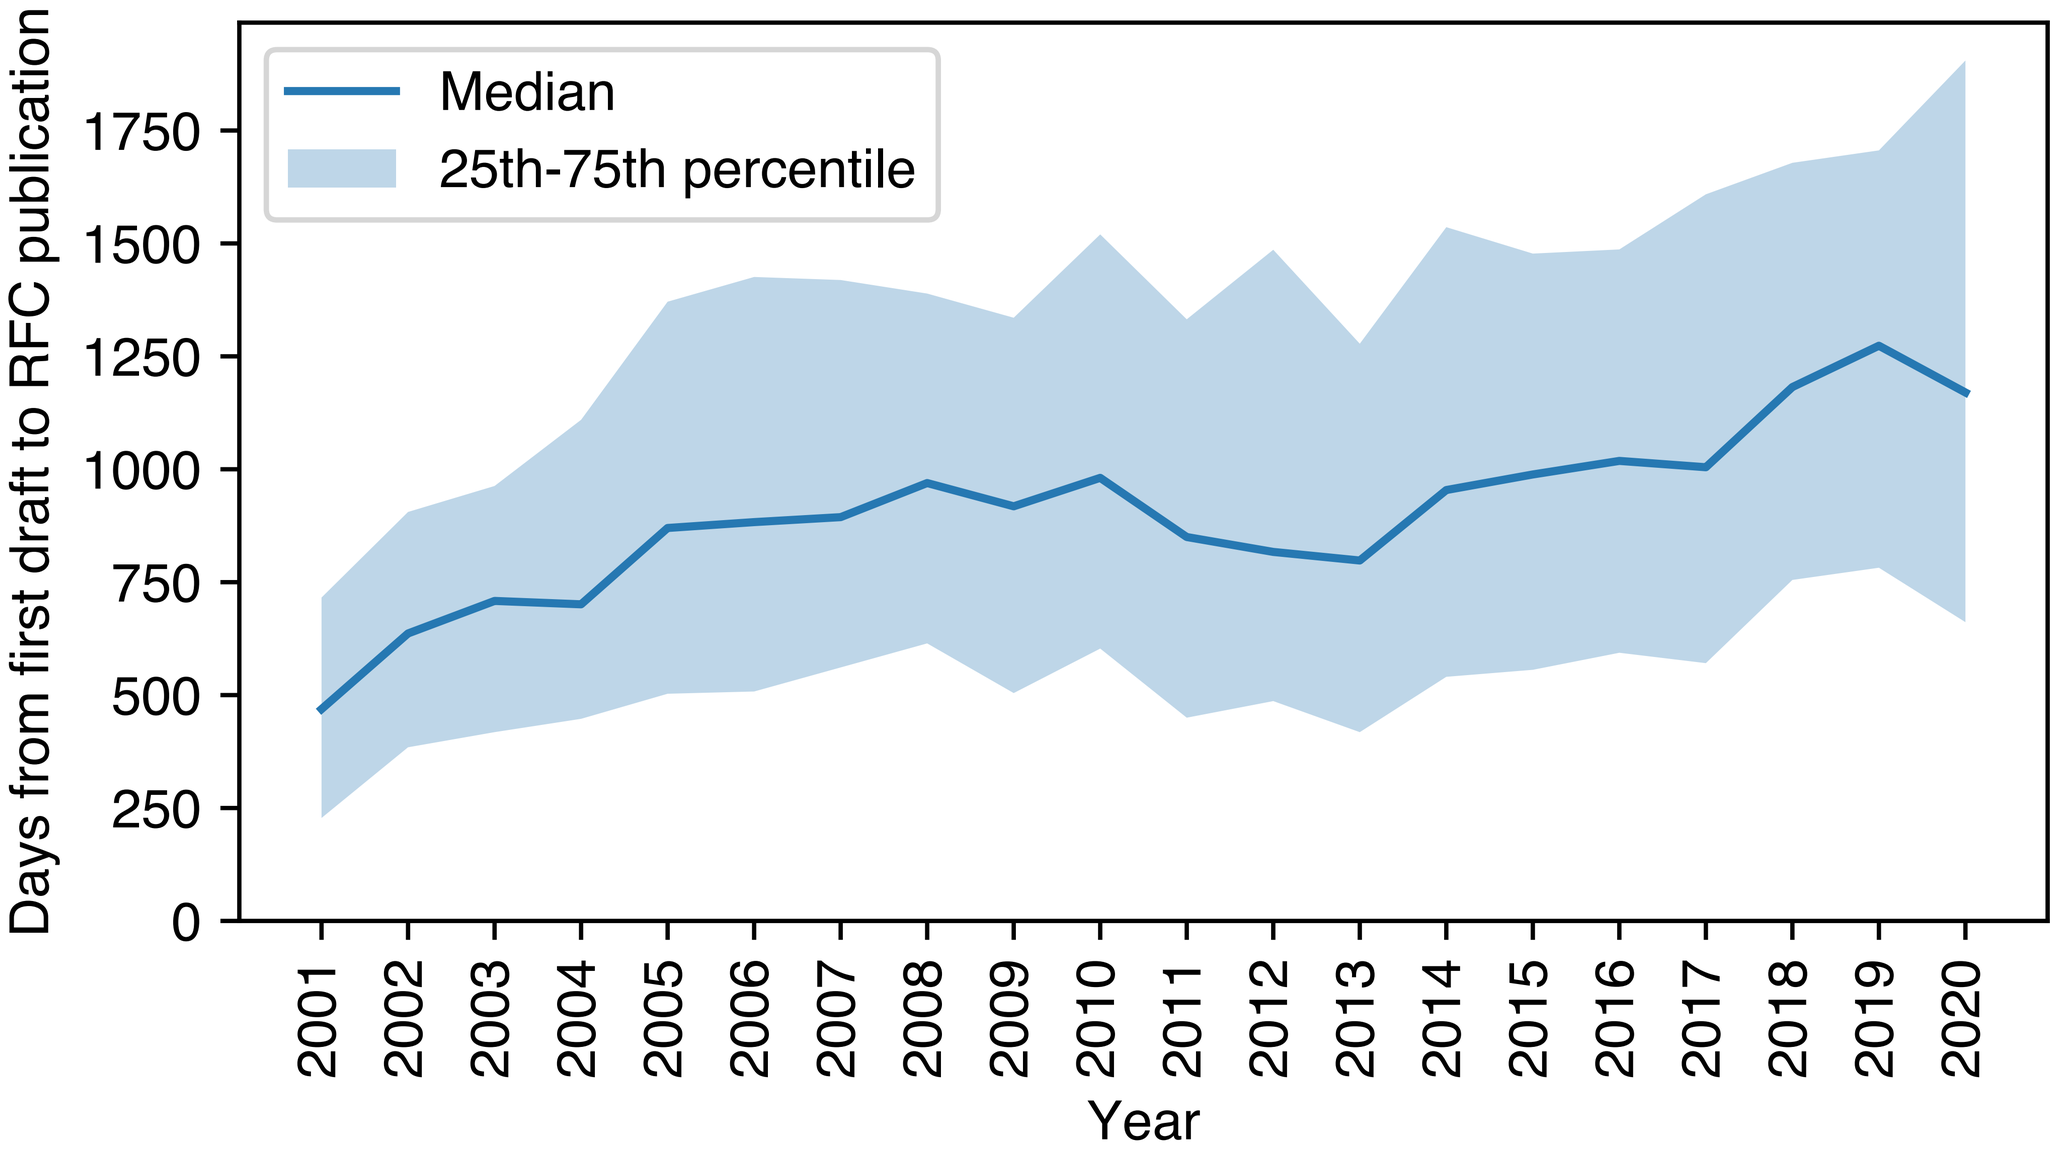
sodestream.github.io



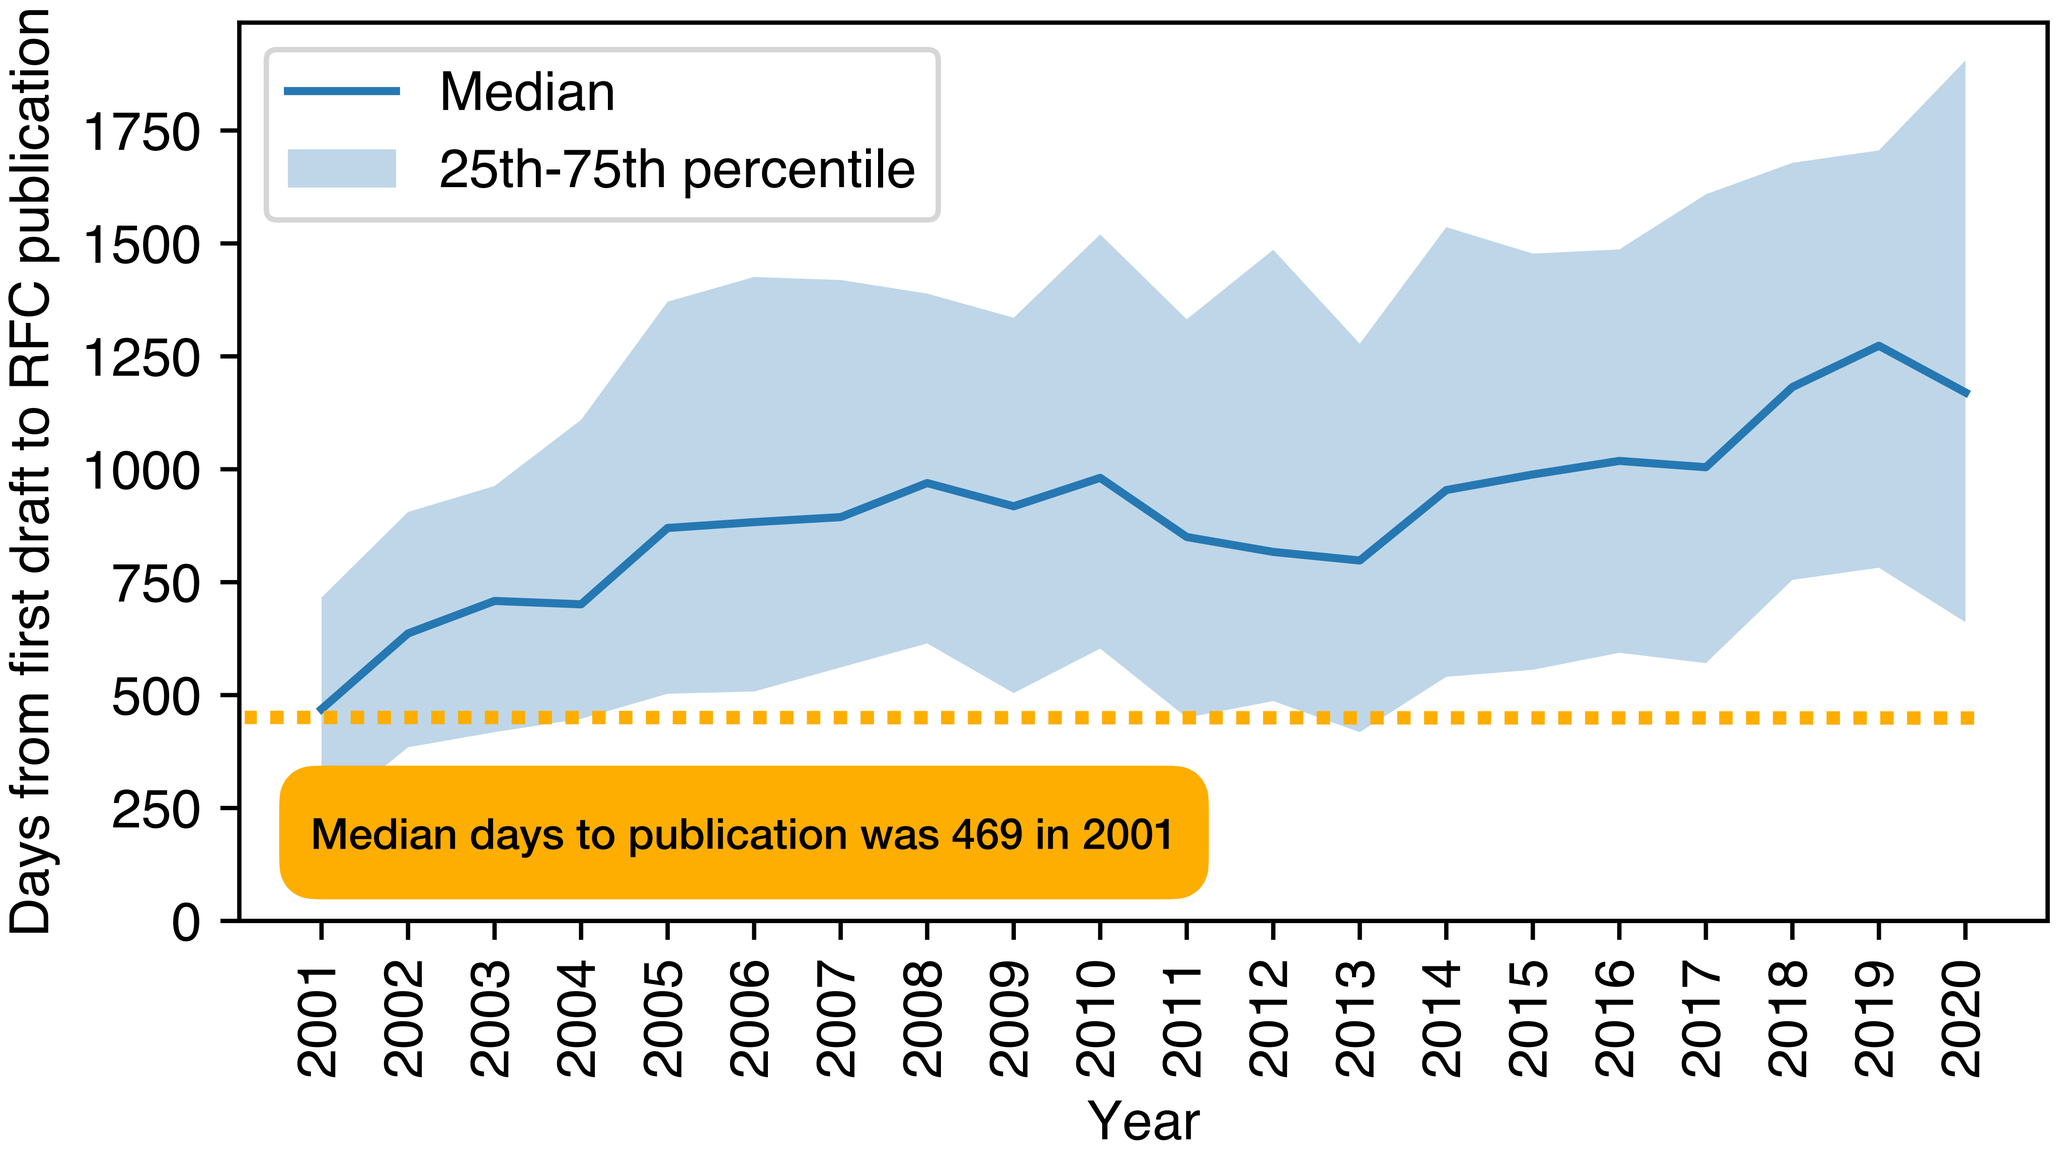
sodestream.github.io



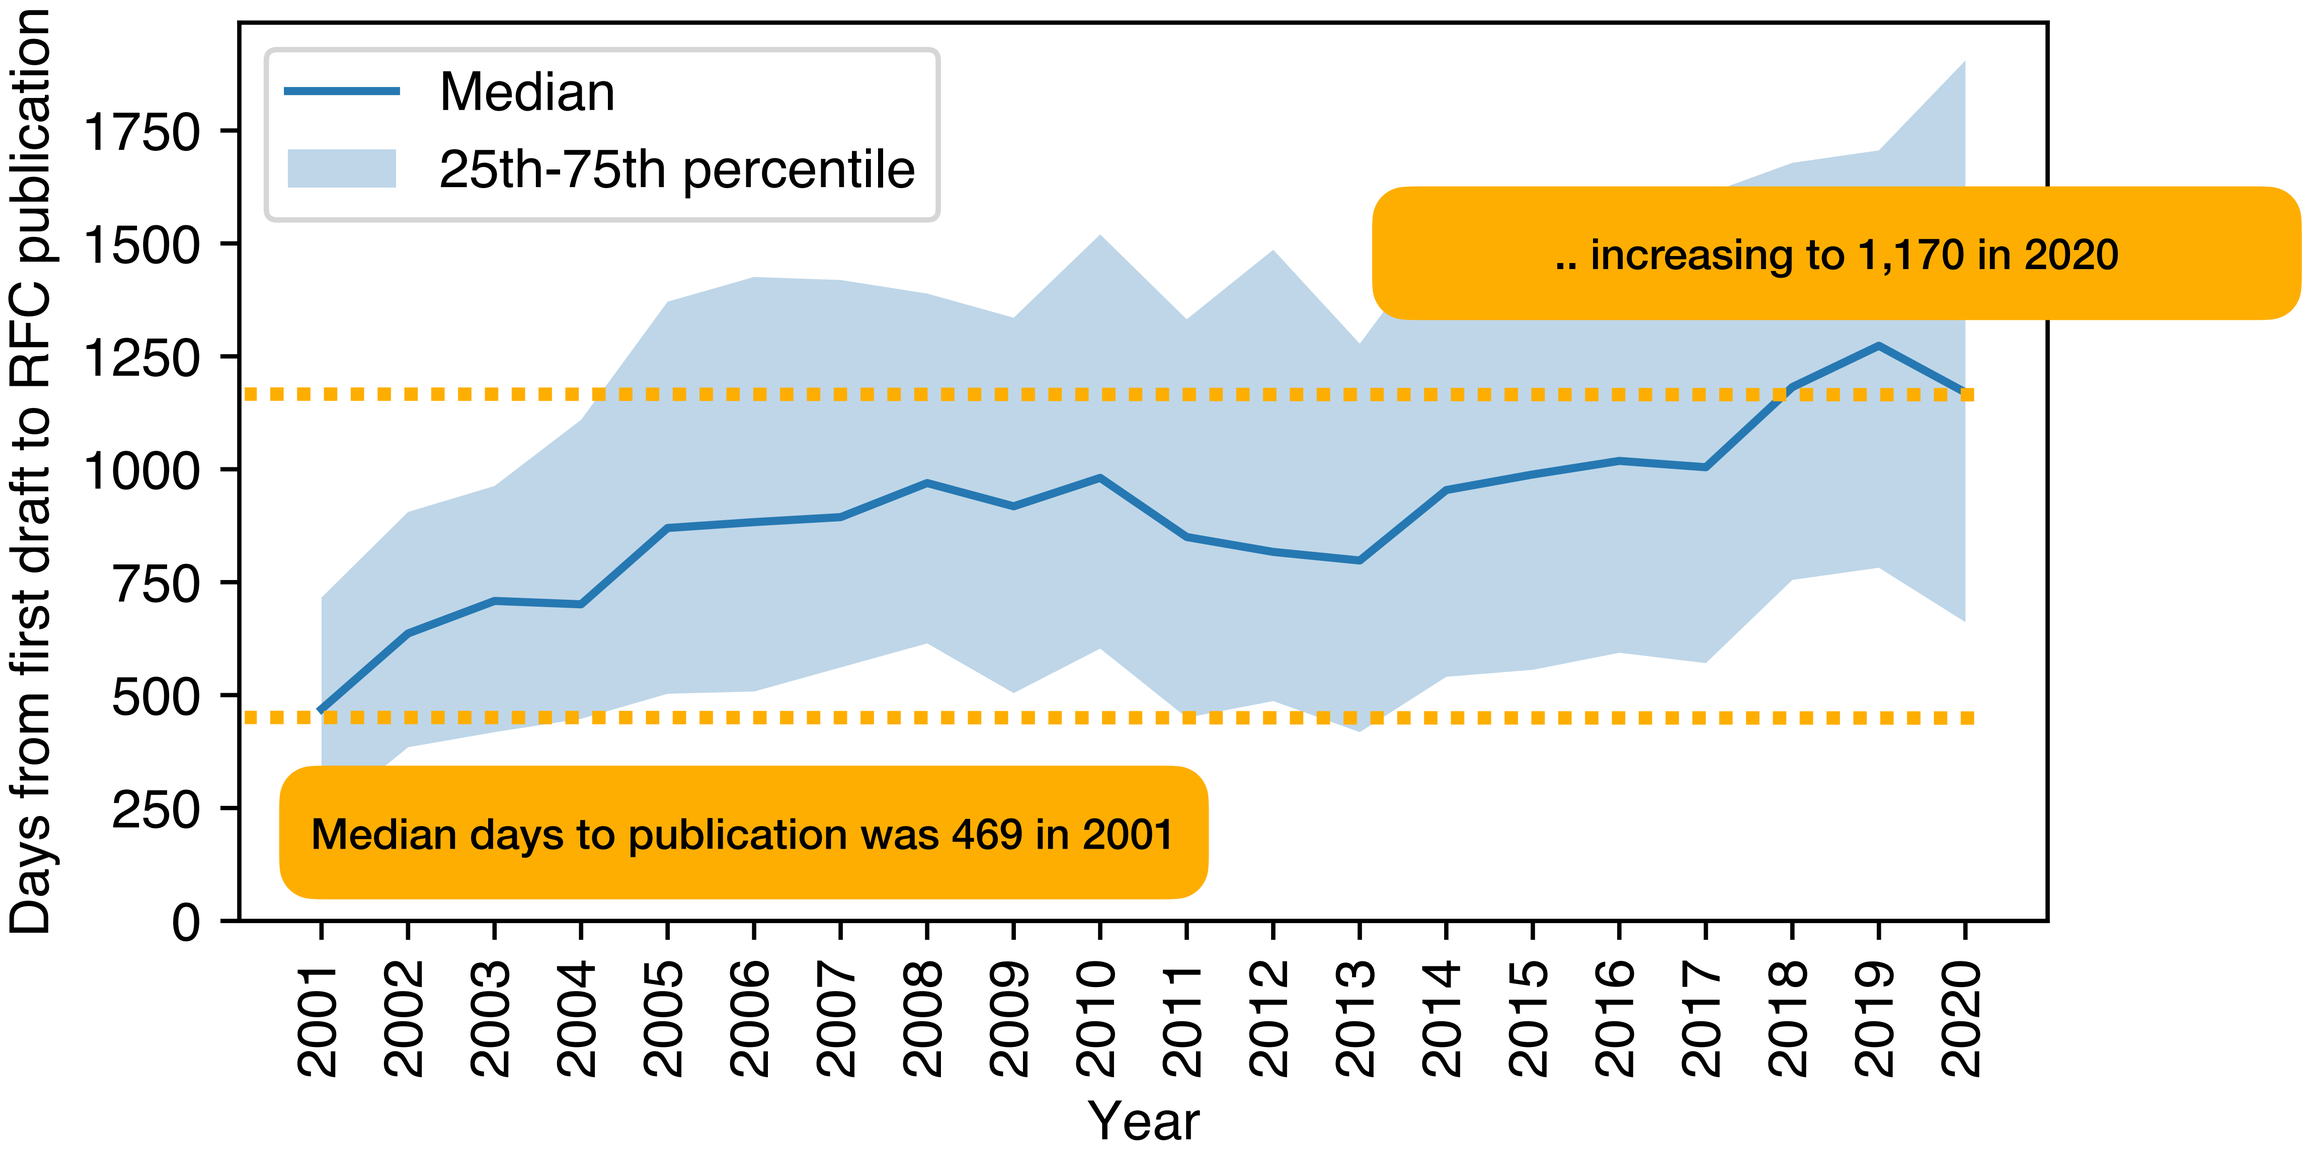
sodestream.github.io



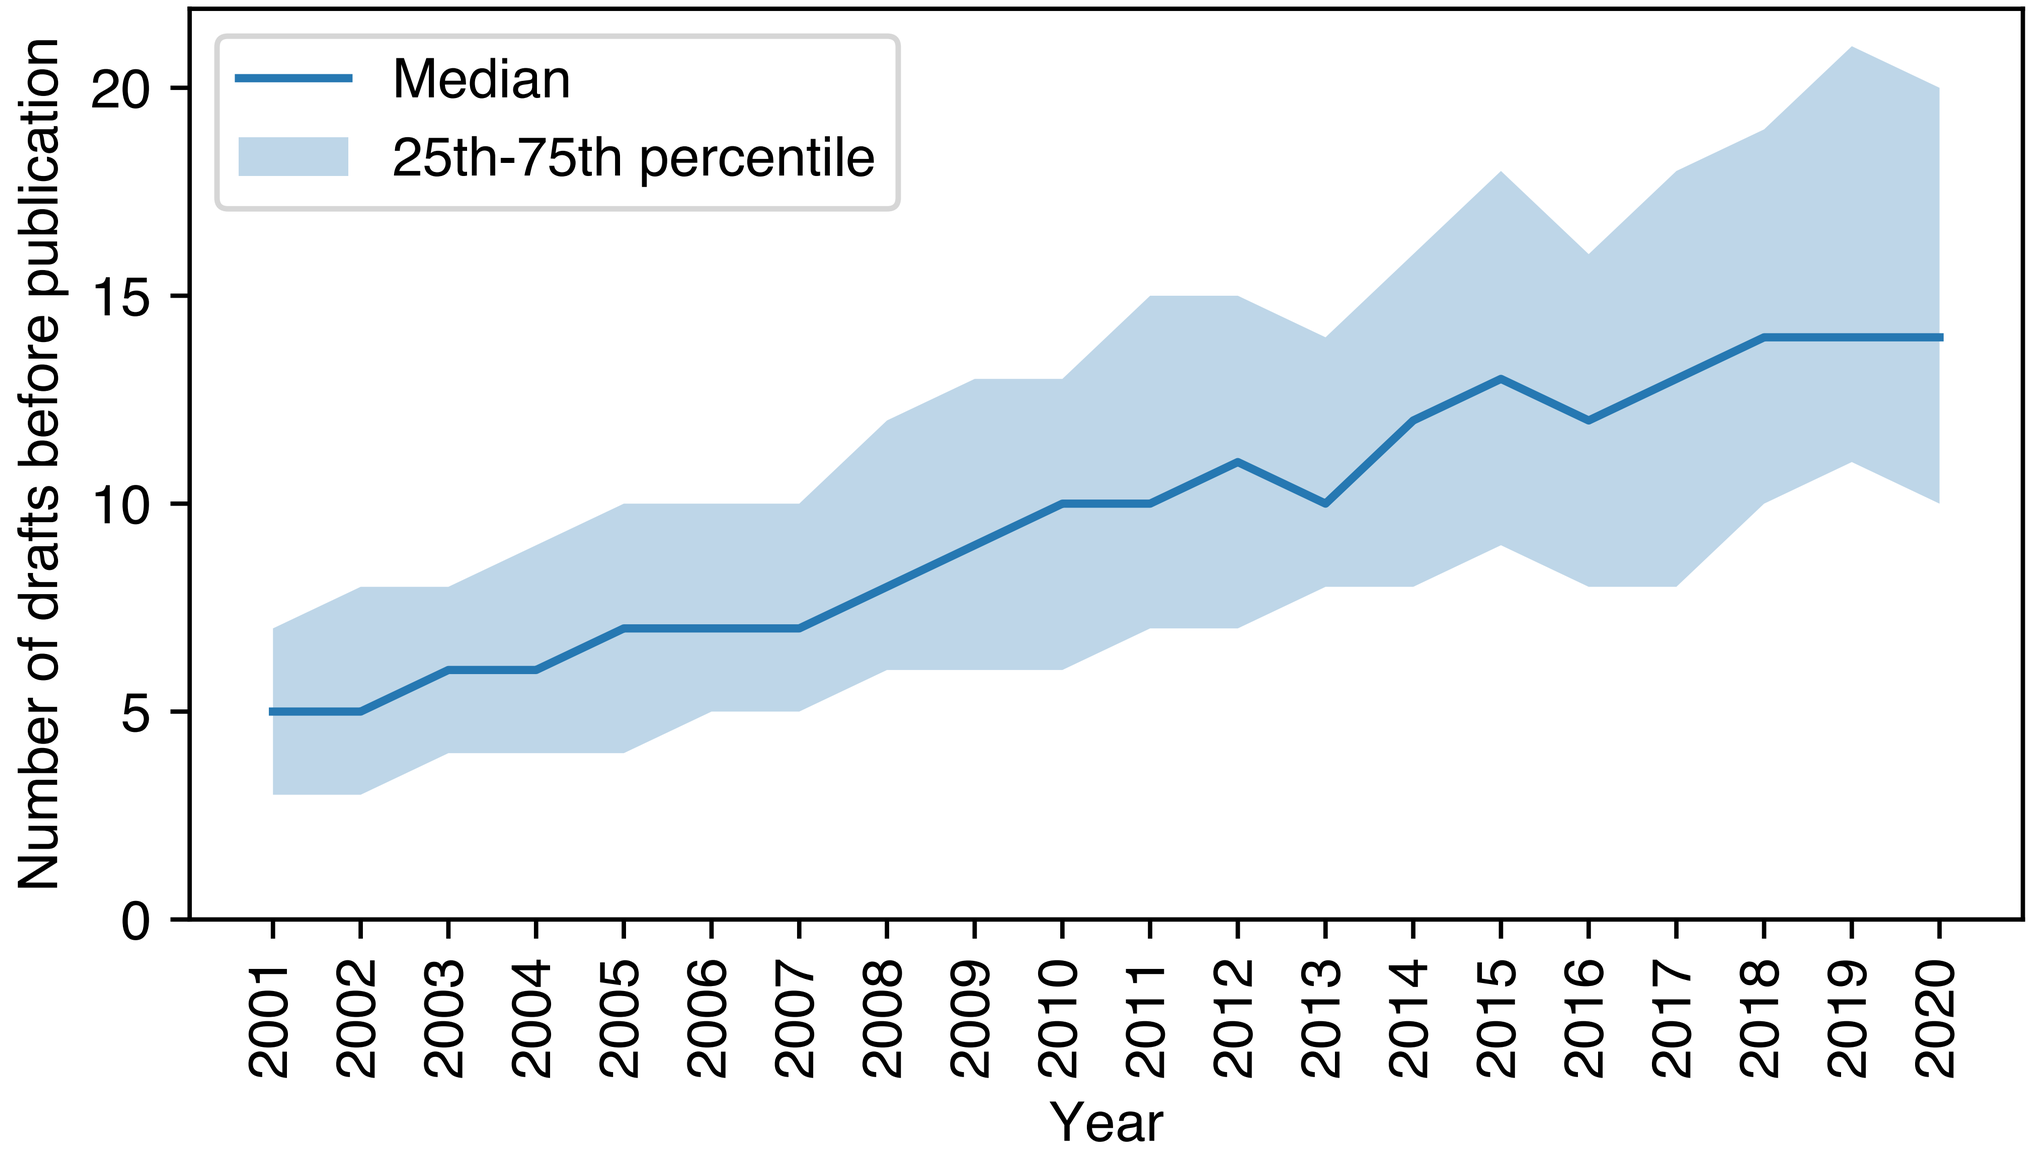
sodestream.github.io



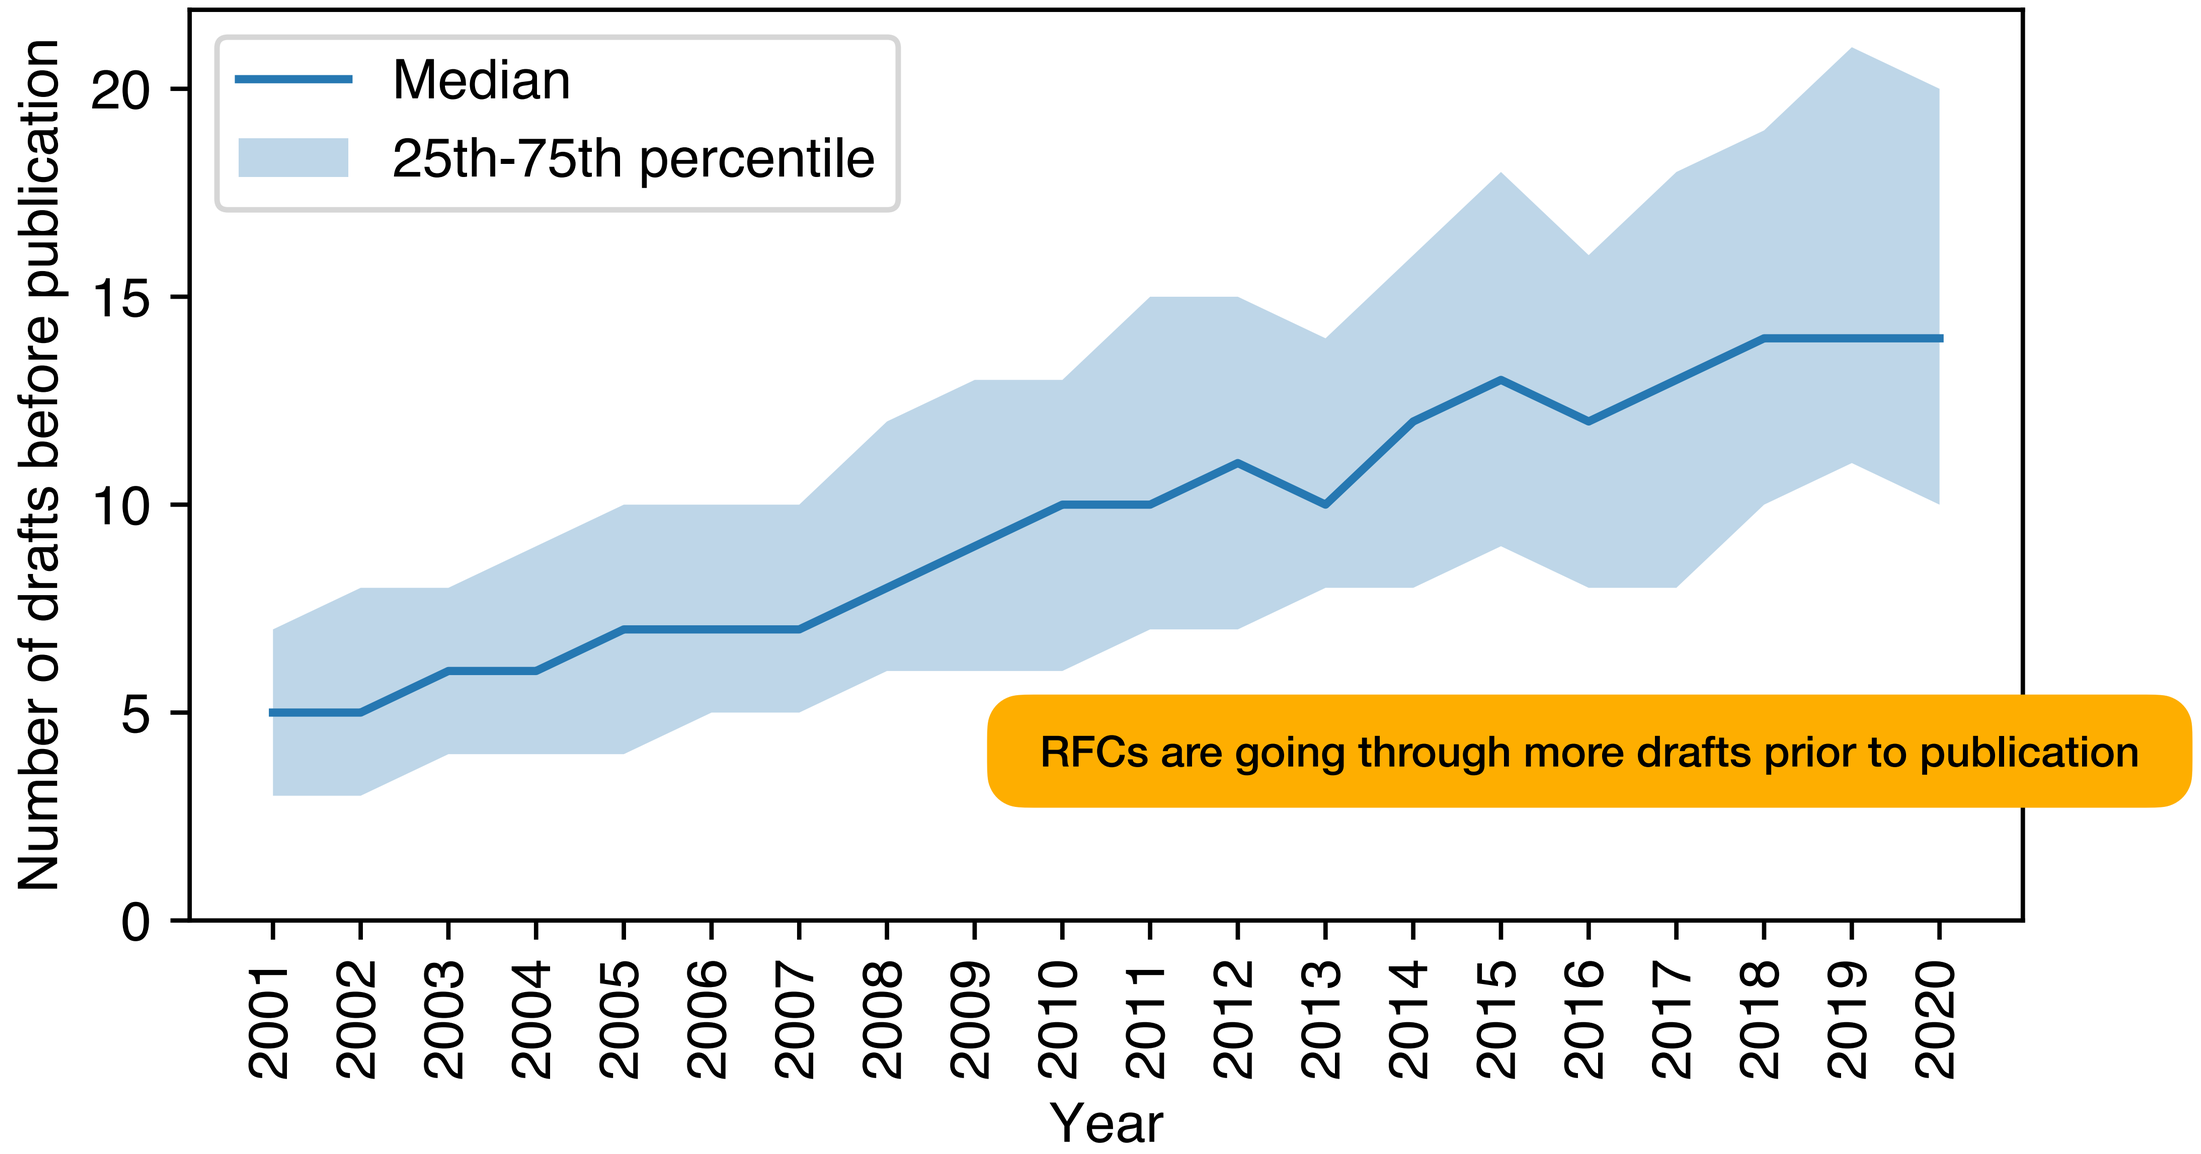
sodestream.github.io



sodestream.github.io



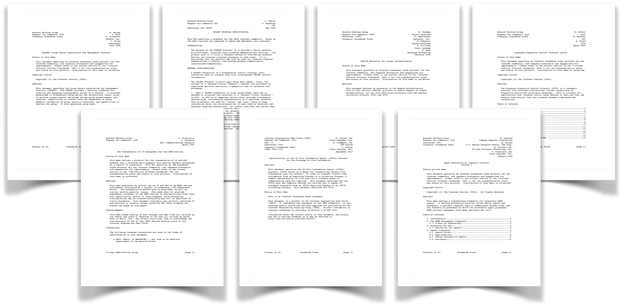
sodestream.github.io



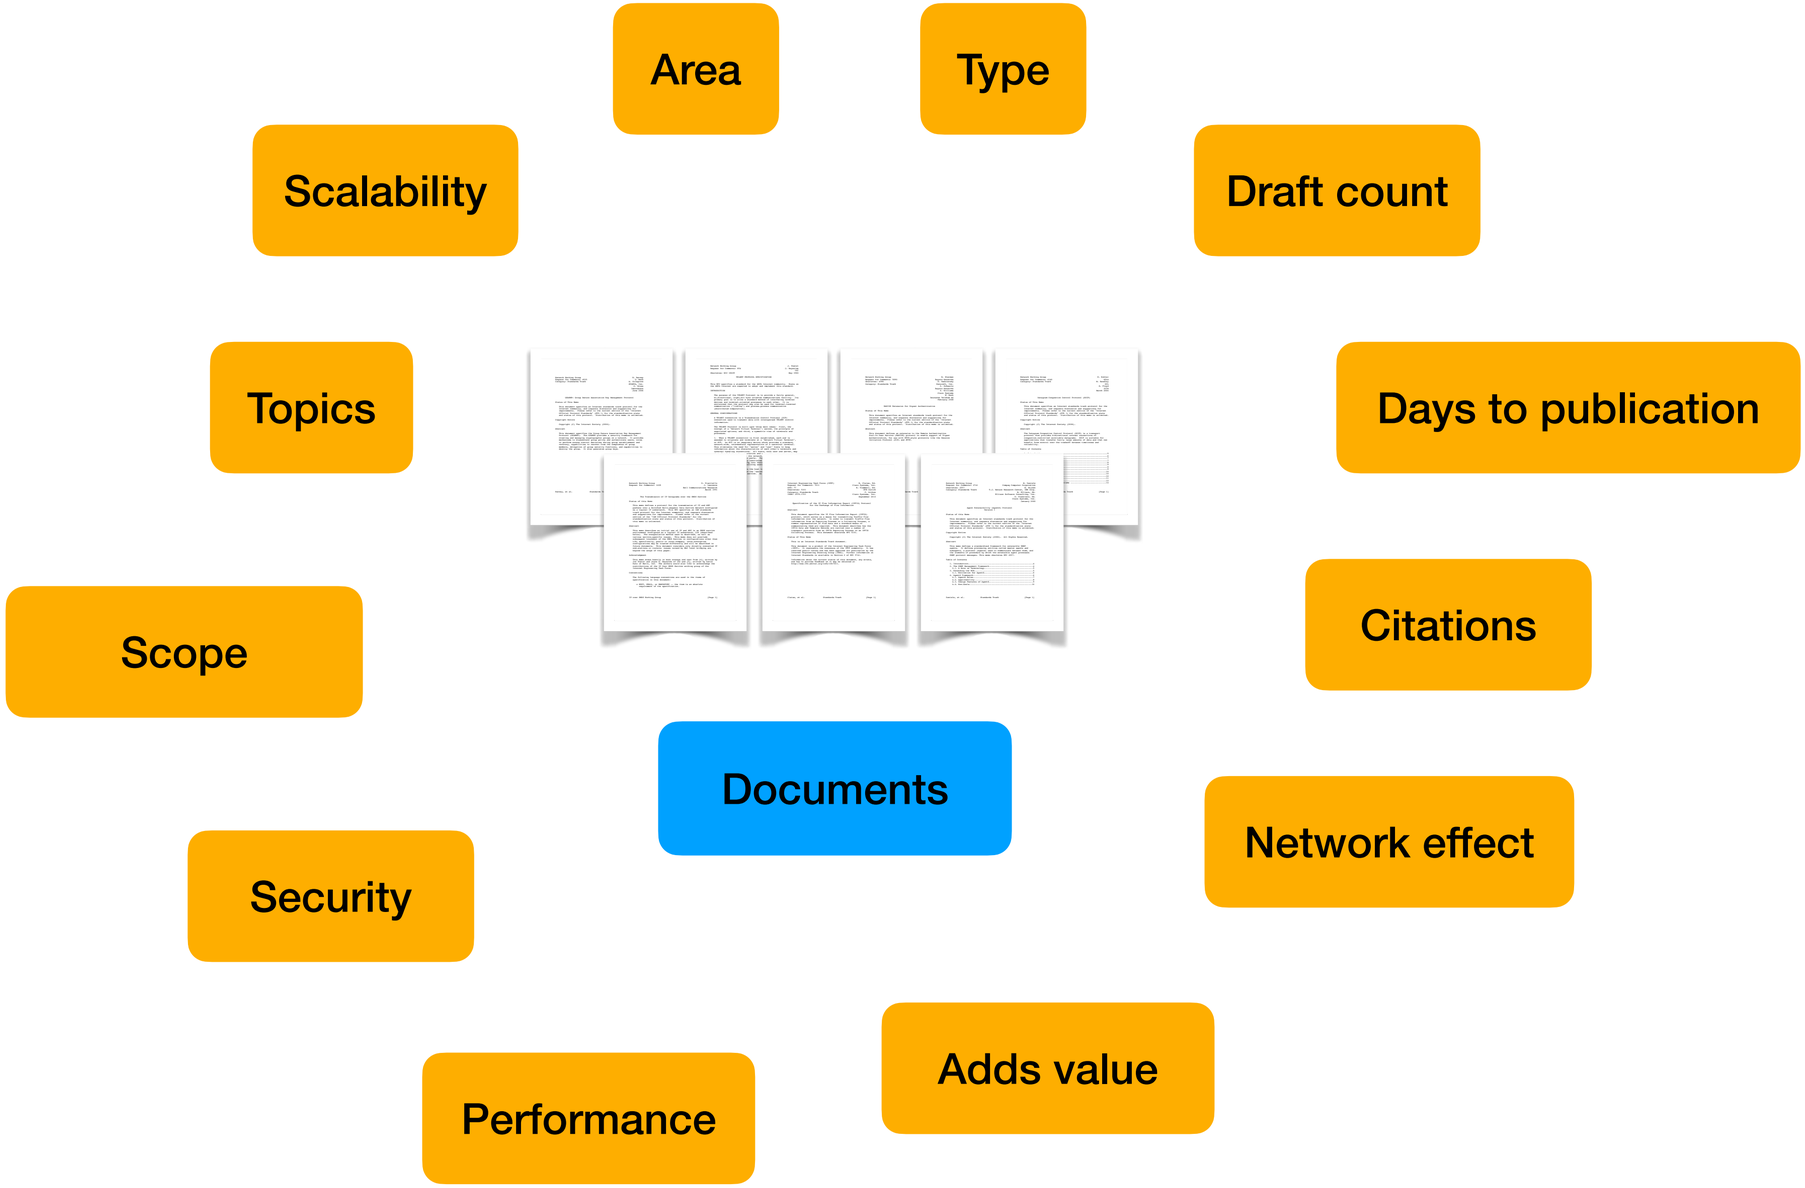
sodestream.github.io

Draft count

Days to publication



Documents



sodestream.github.io

Nikkhah et al., "*A Statistical Exploration of Protocol Adoption*". IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642

Area

Type

Scalability

Draft count

Topics



Days to publication

Other document-based features described in the paper

Scope

Citations

Documents

Network effect

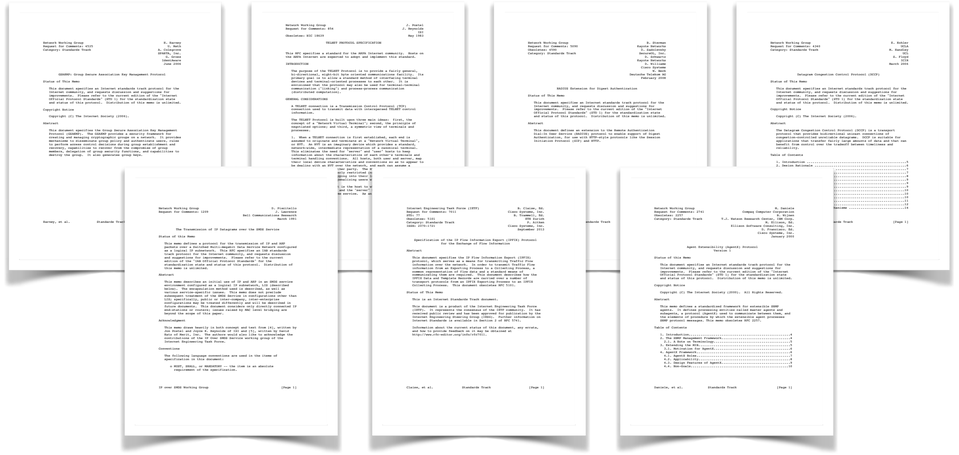
Security

Adds value

Performance

sodestream.github.io

Nikkhah et al., “*A Statistical Exploration of Protocol Adoption*”.
IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642



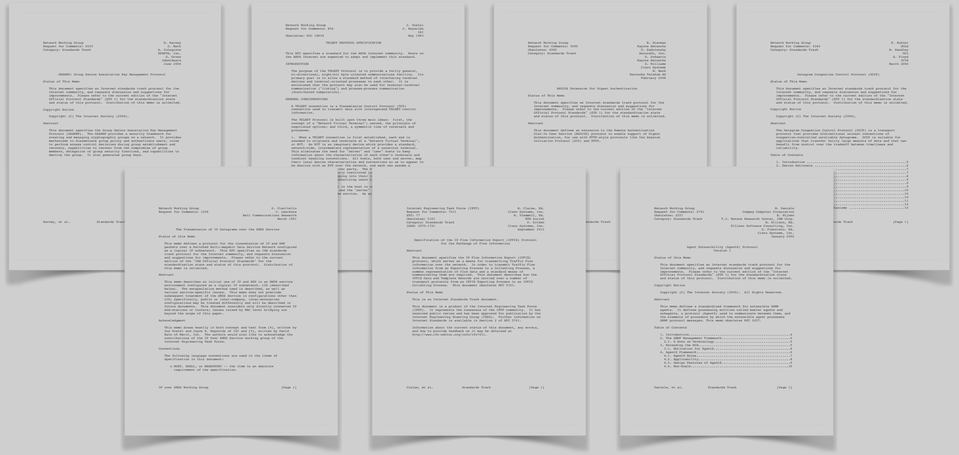
Documents

Authors

IETF®

Mailing lists

sodestream.github.io



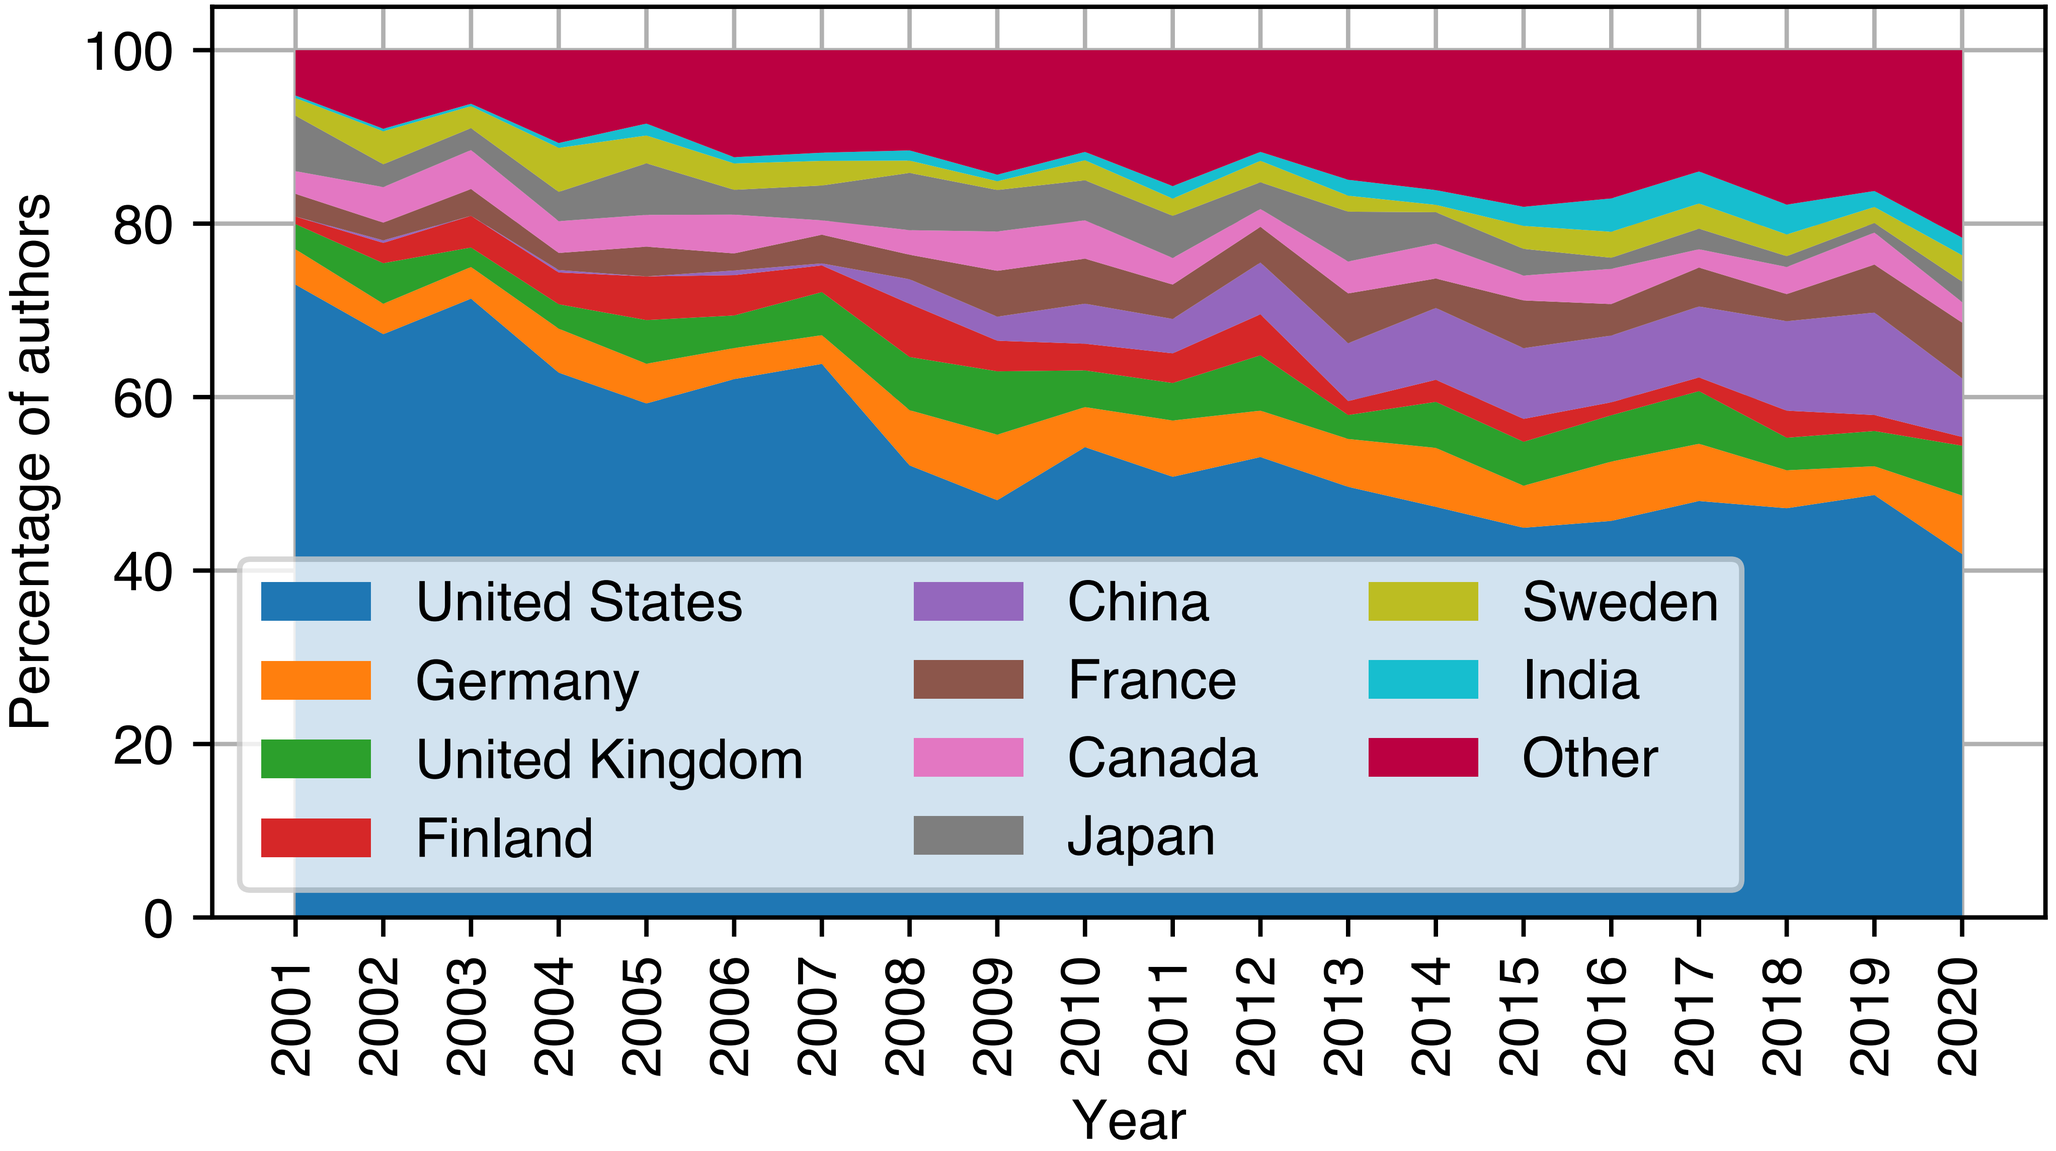
Documents

Authors

IETF®

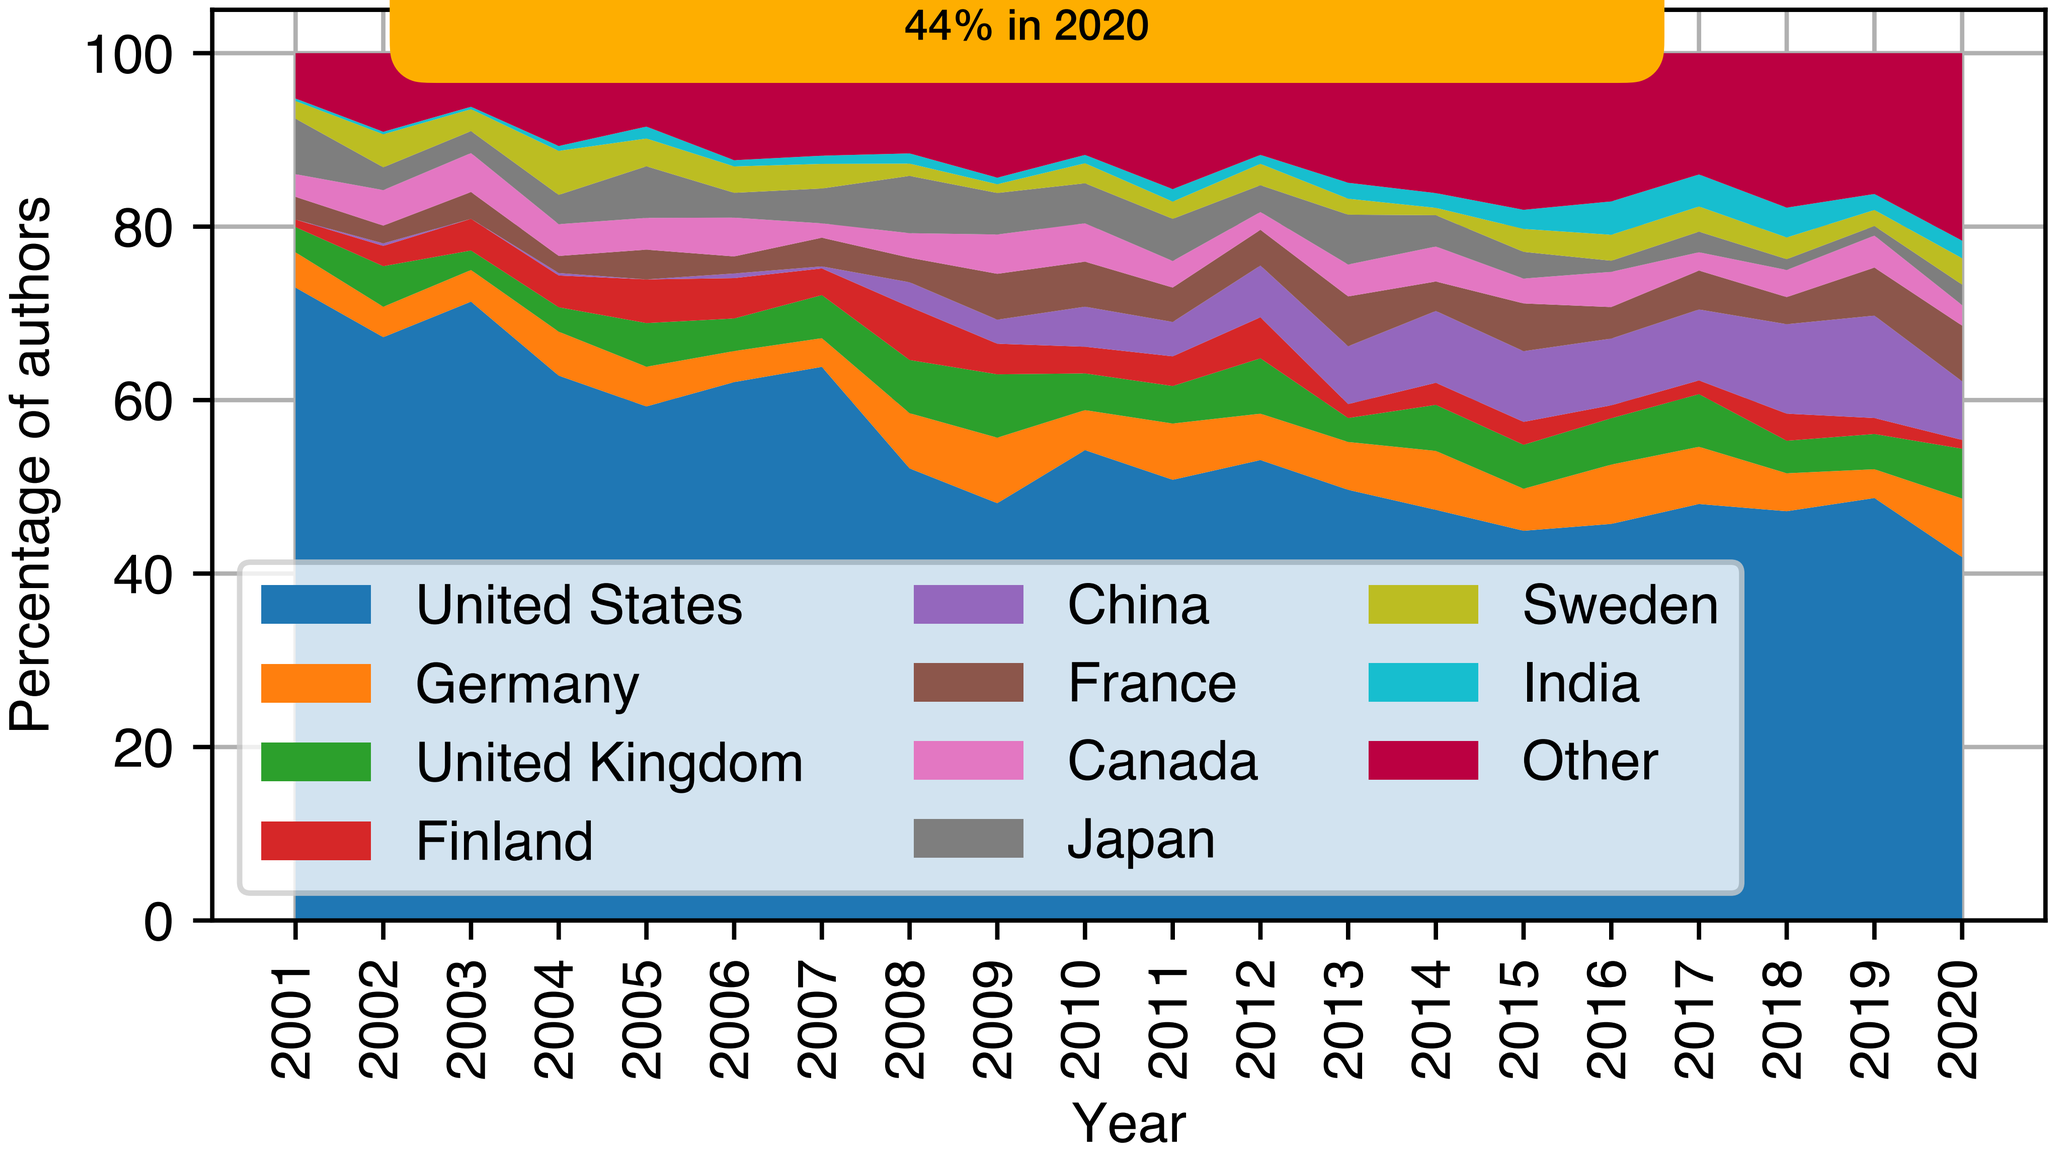
Mailing lists

sodestream.github.io

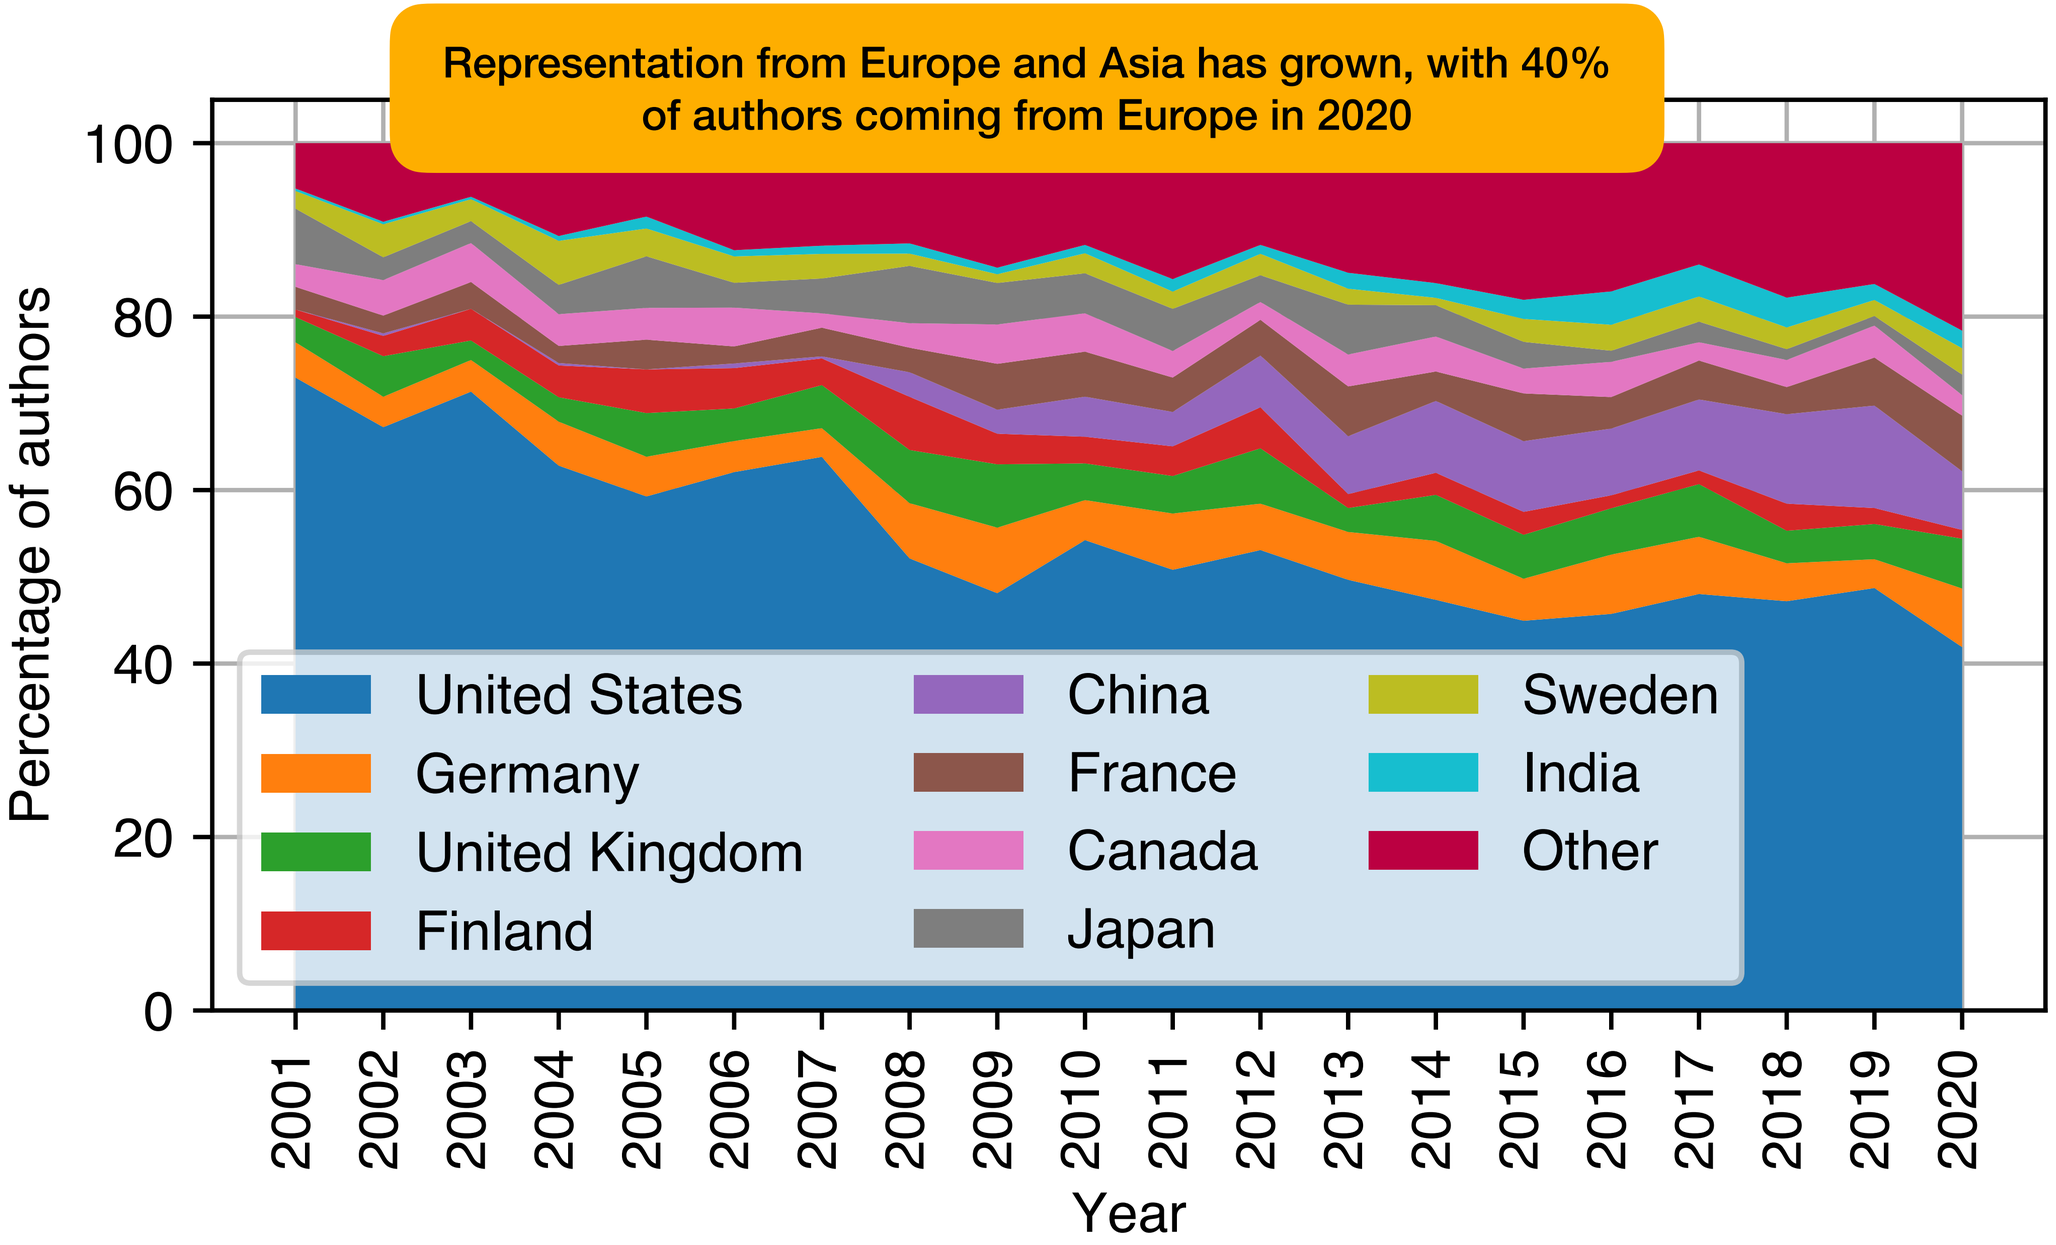


sodestream.github.io

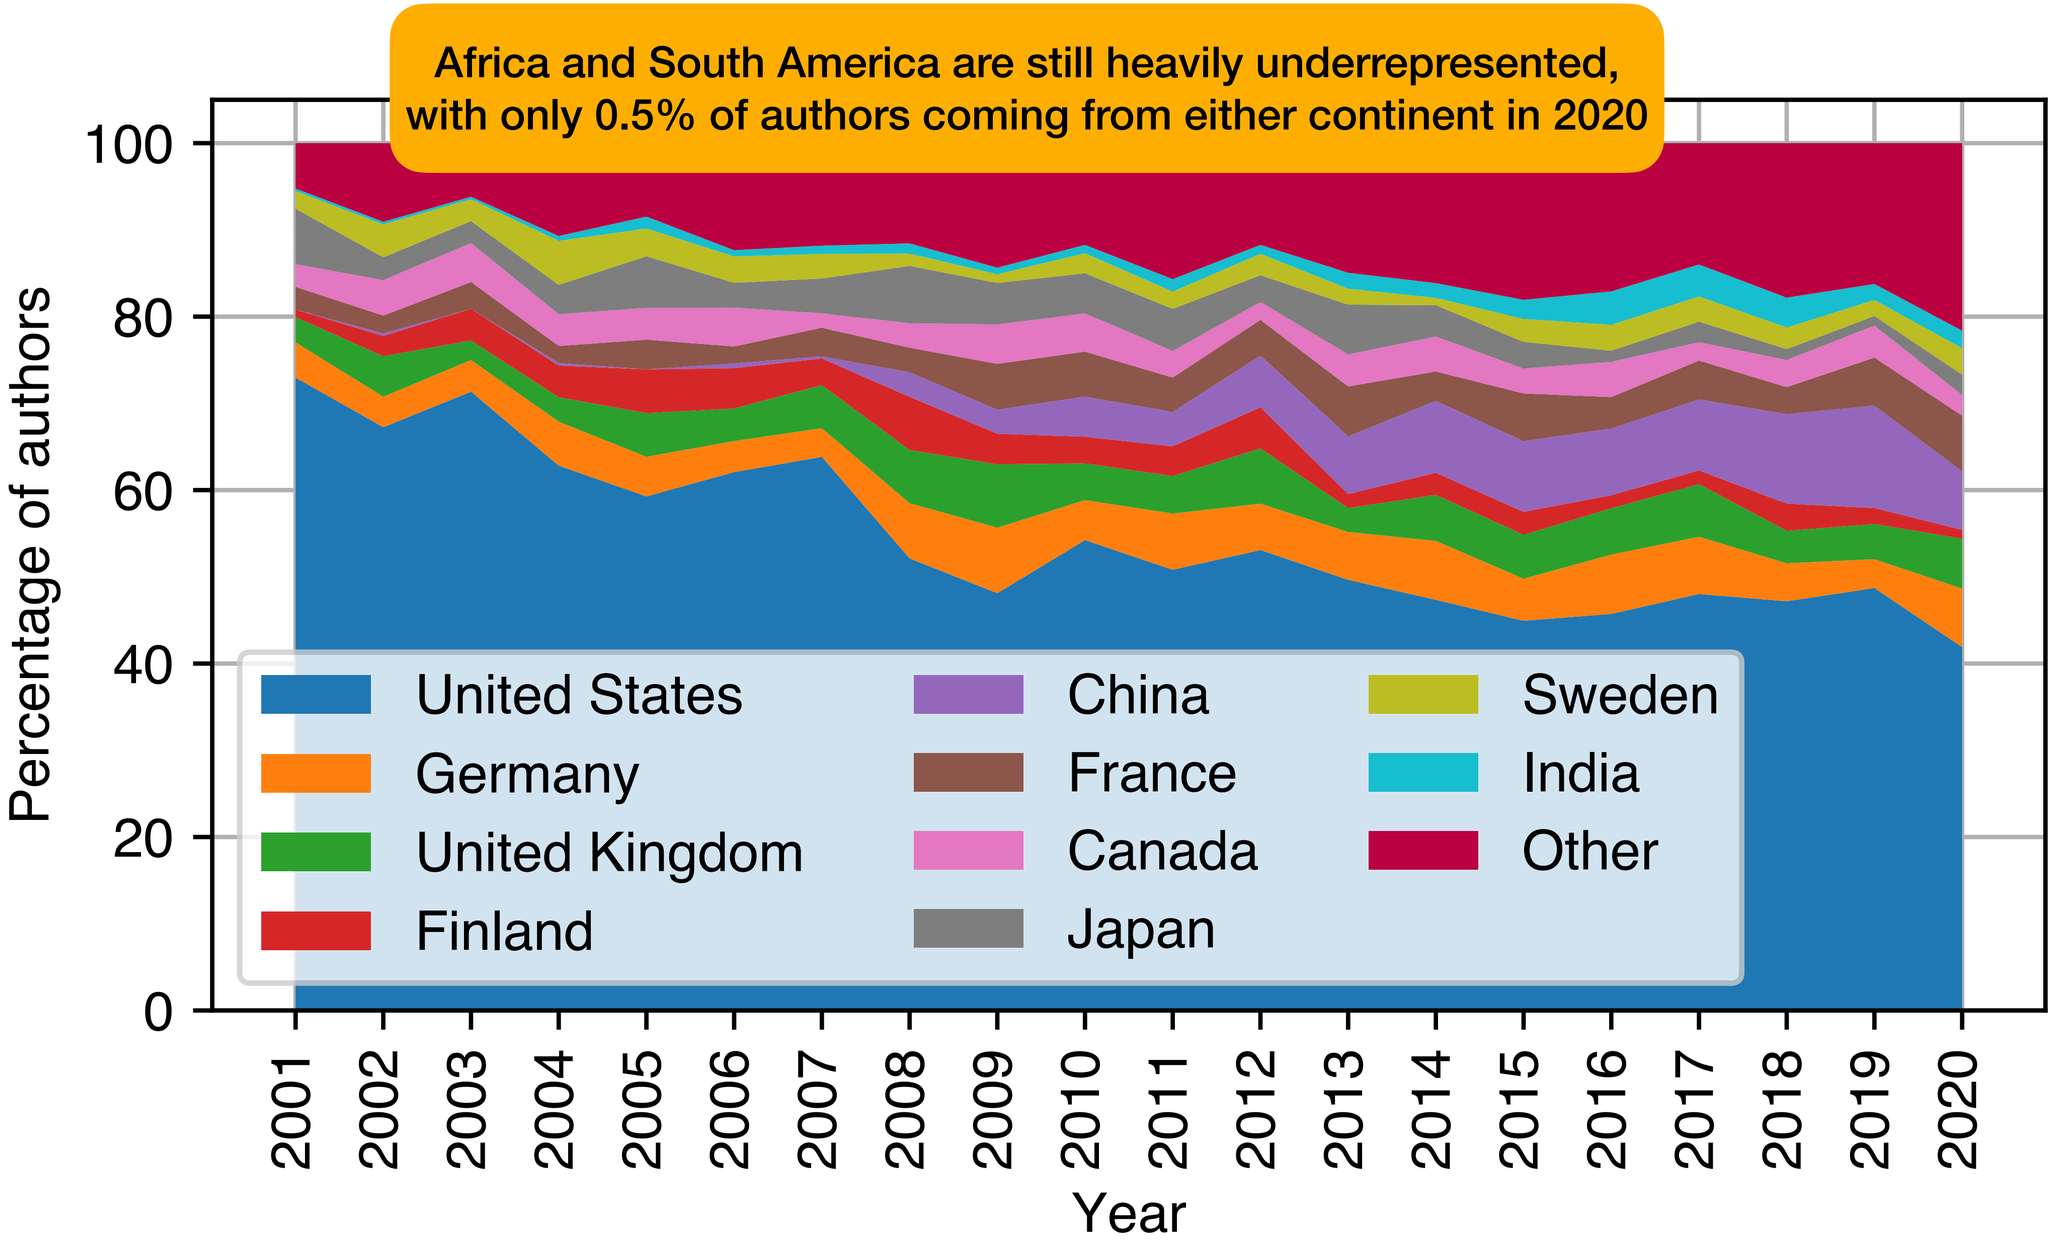
75% of authors were from North America in 2001, falling to 44% in 2020



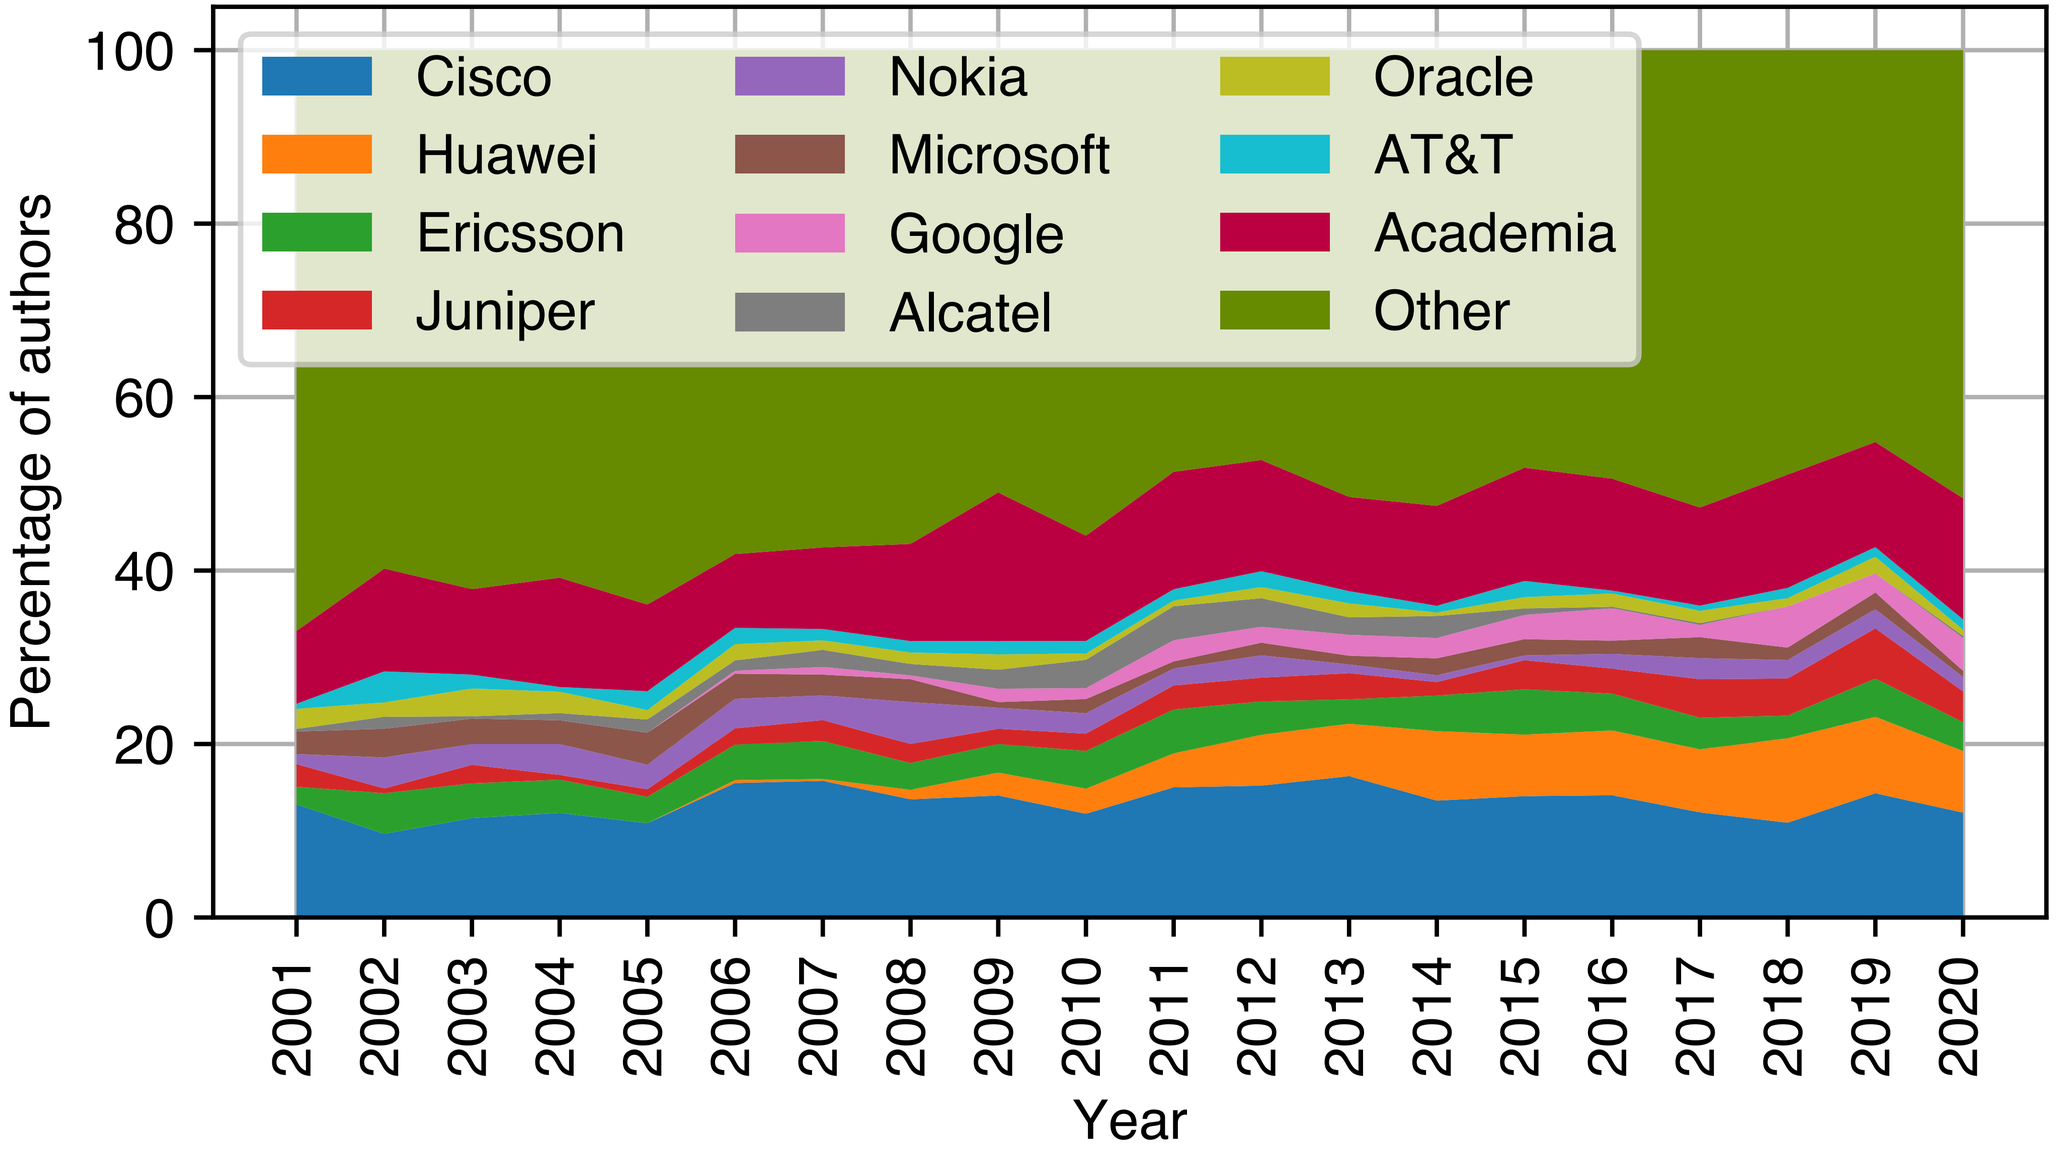
sodestream.github.io



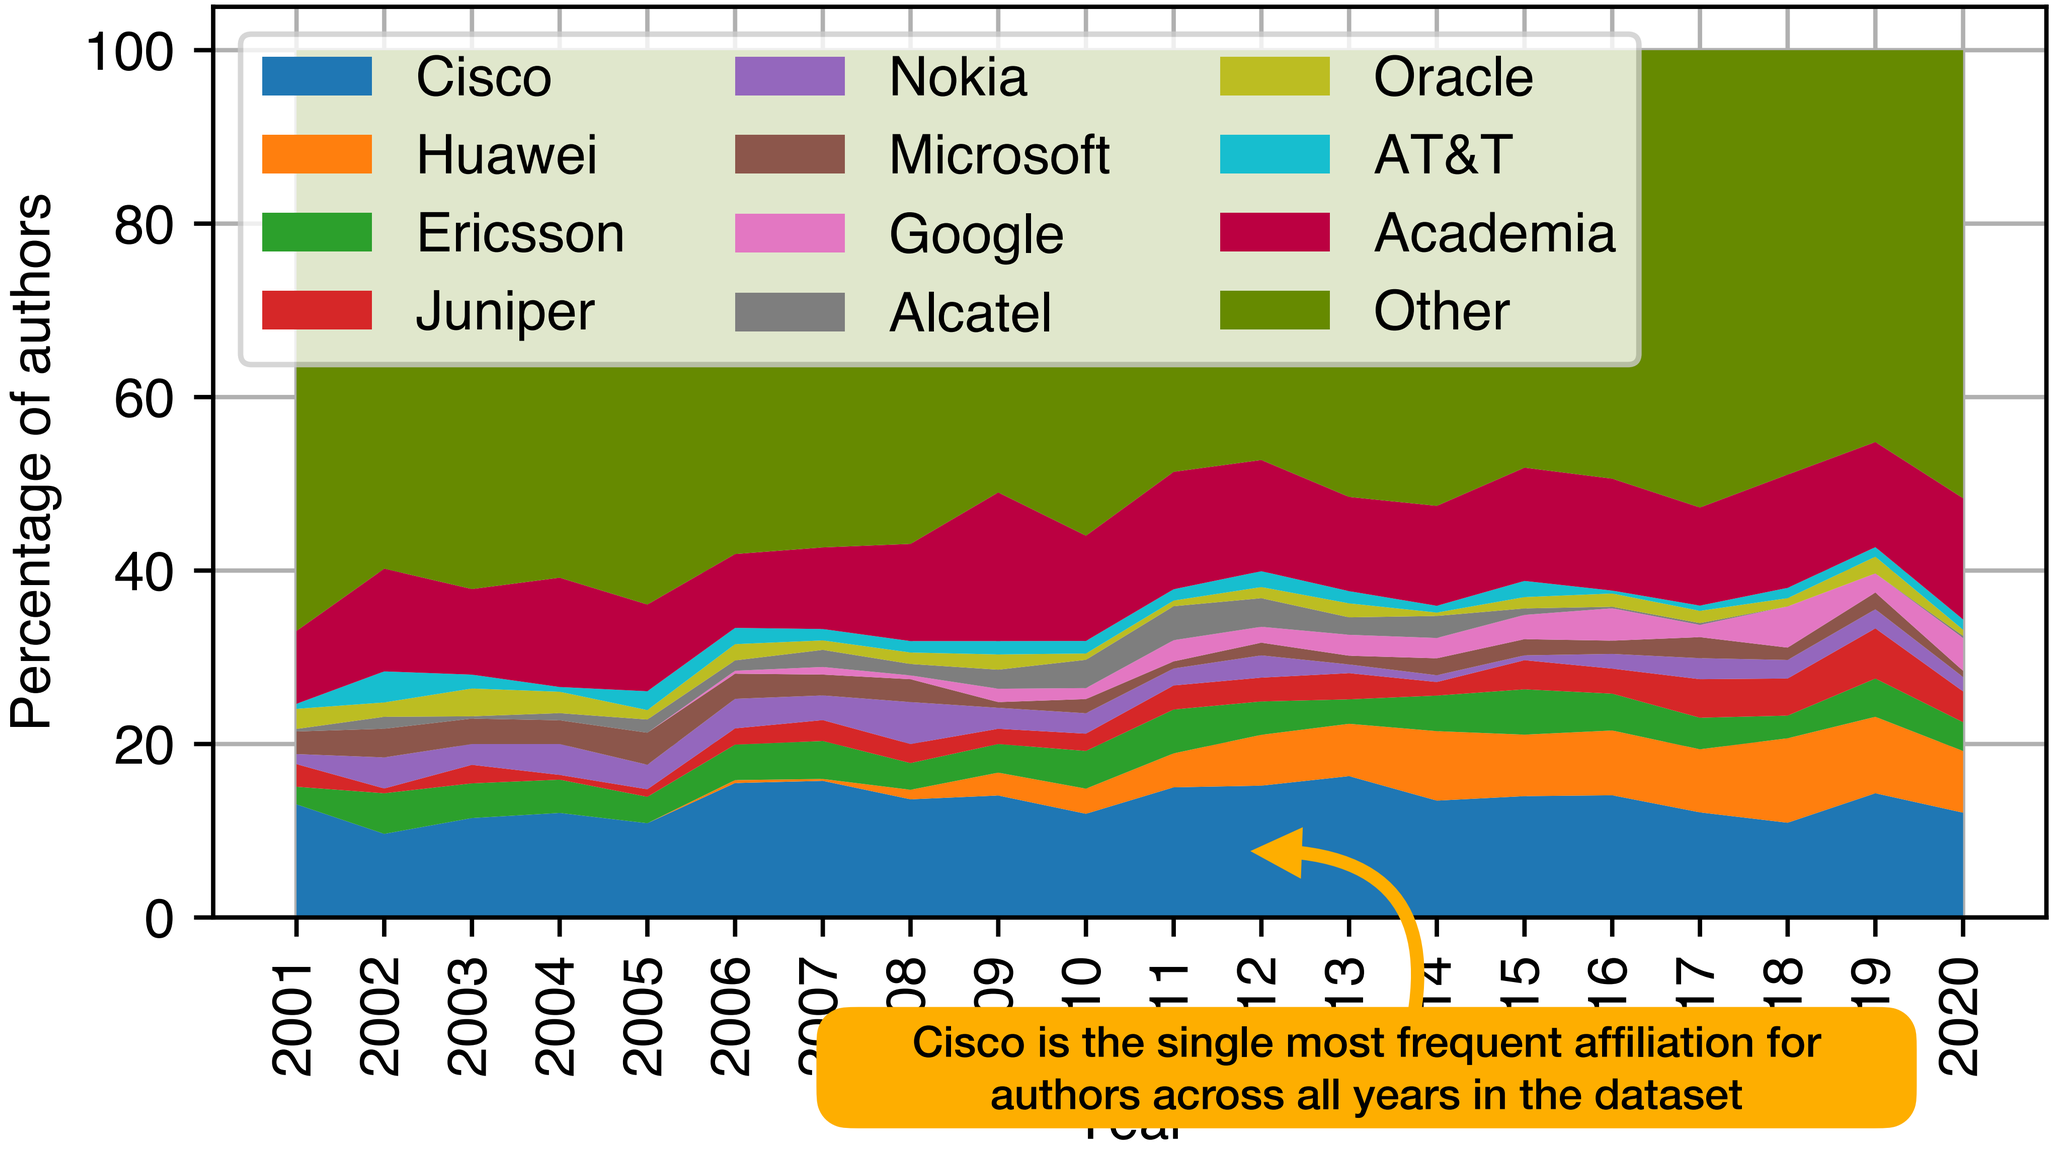
sodestream.github.io



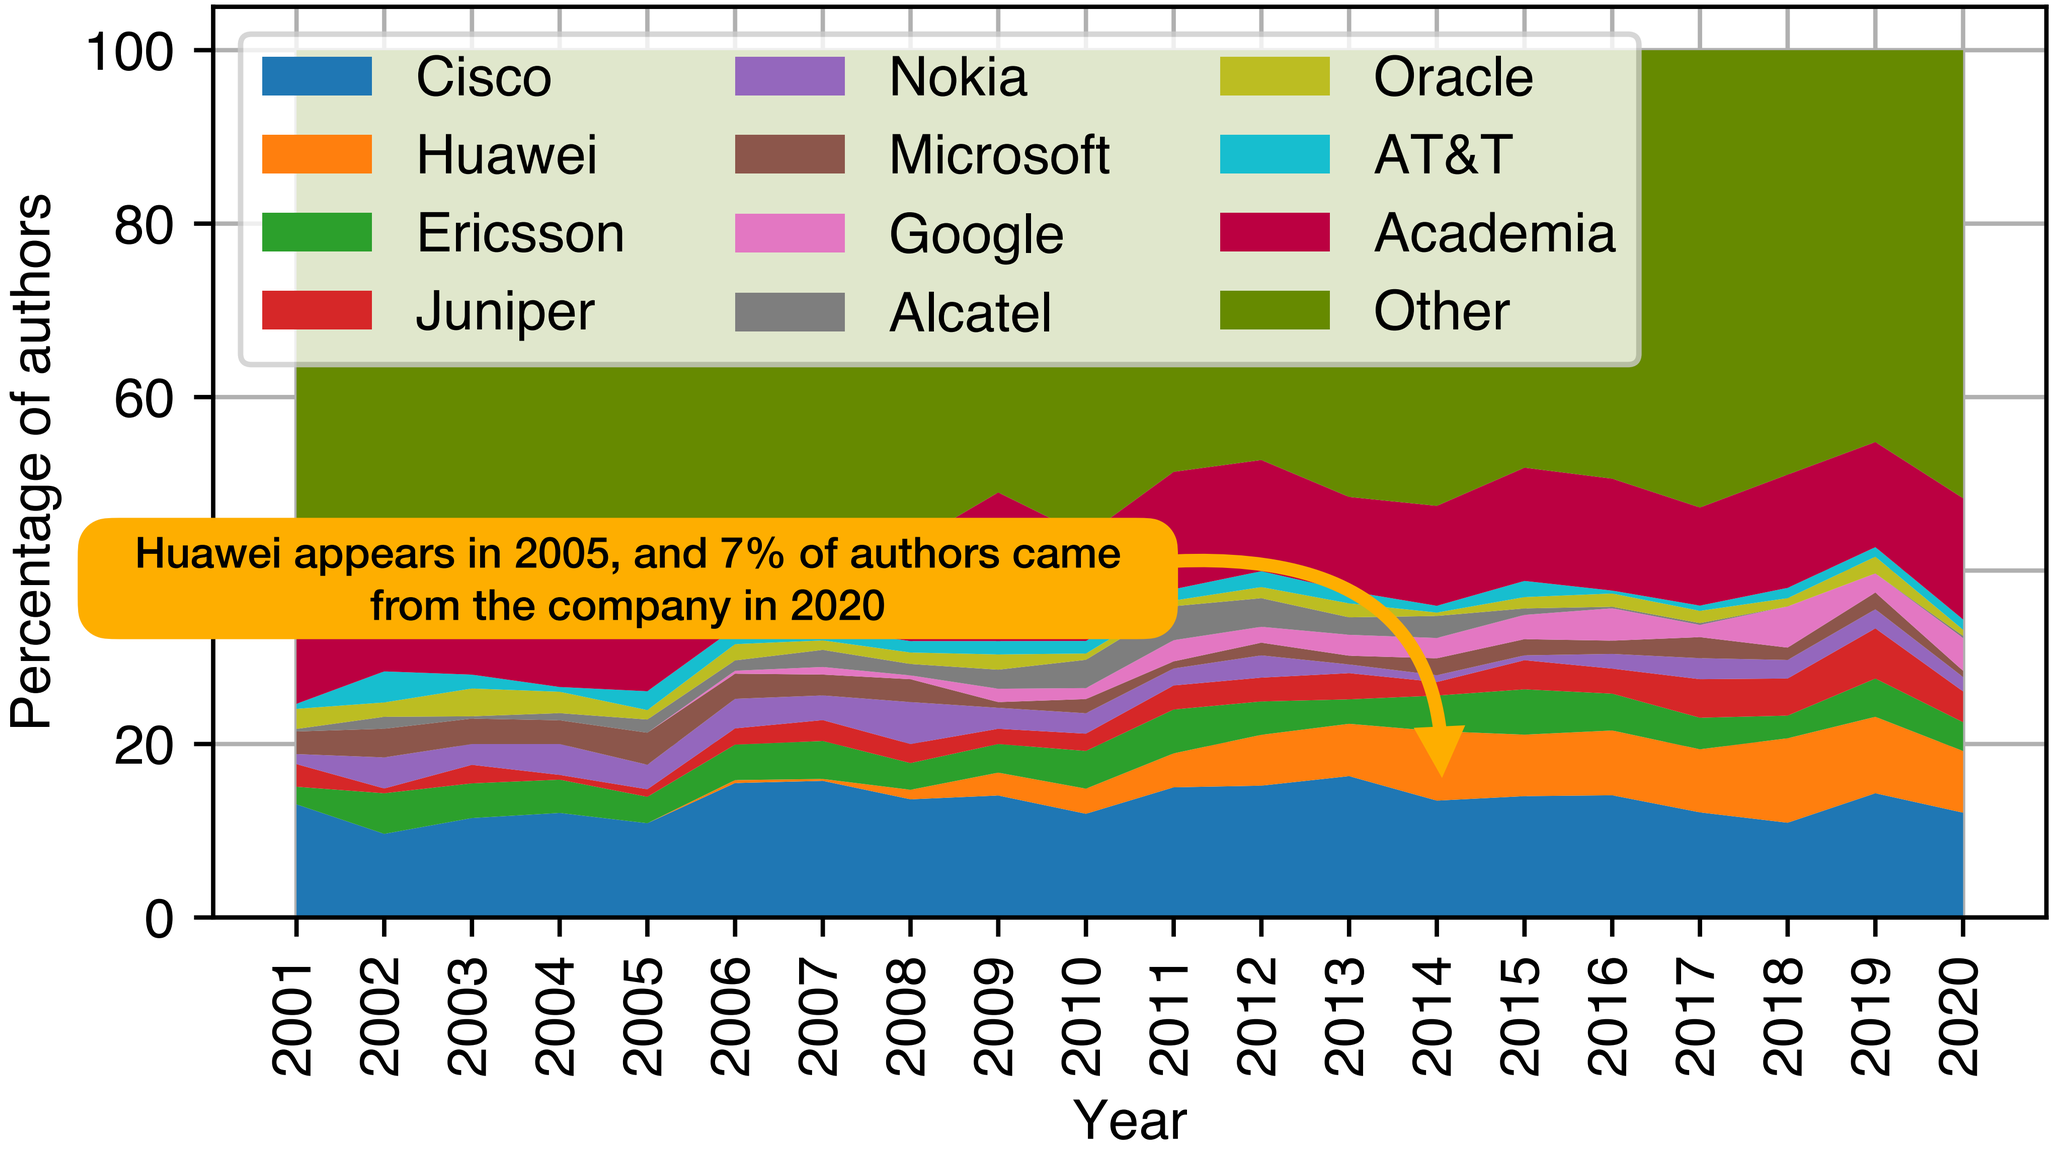
sodestream.github.io



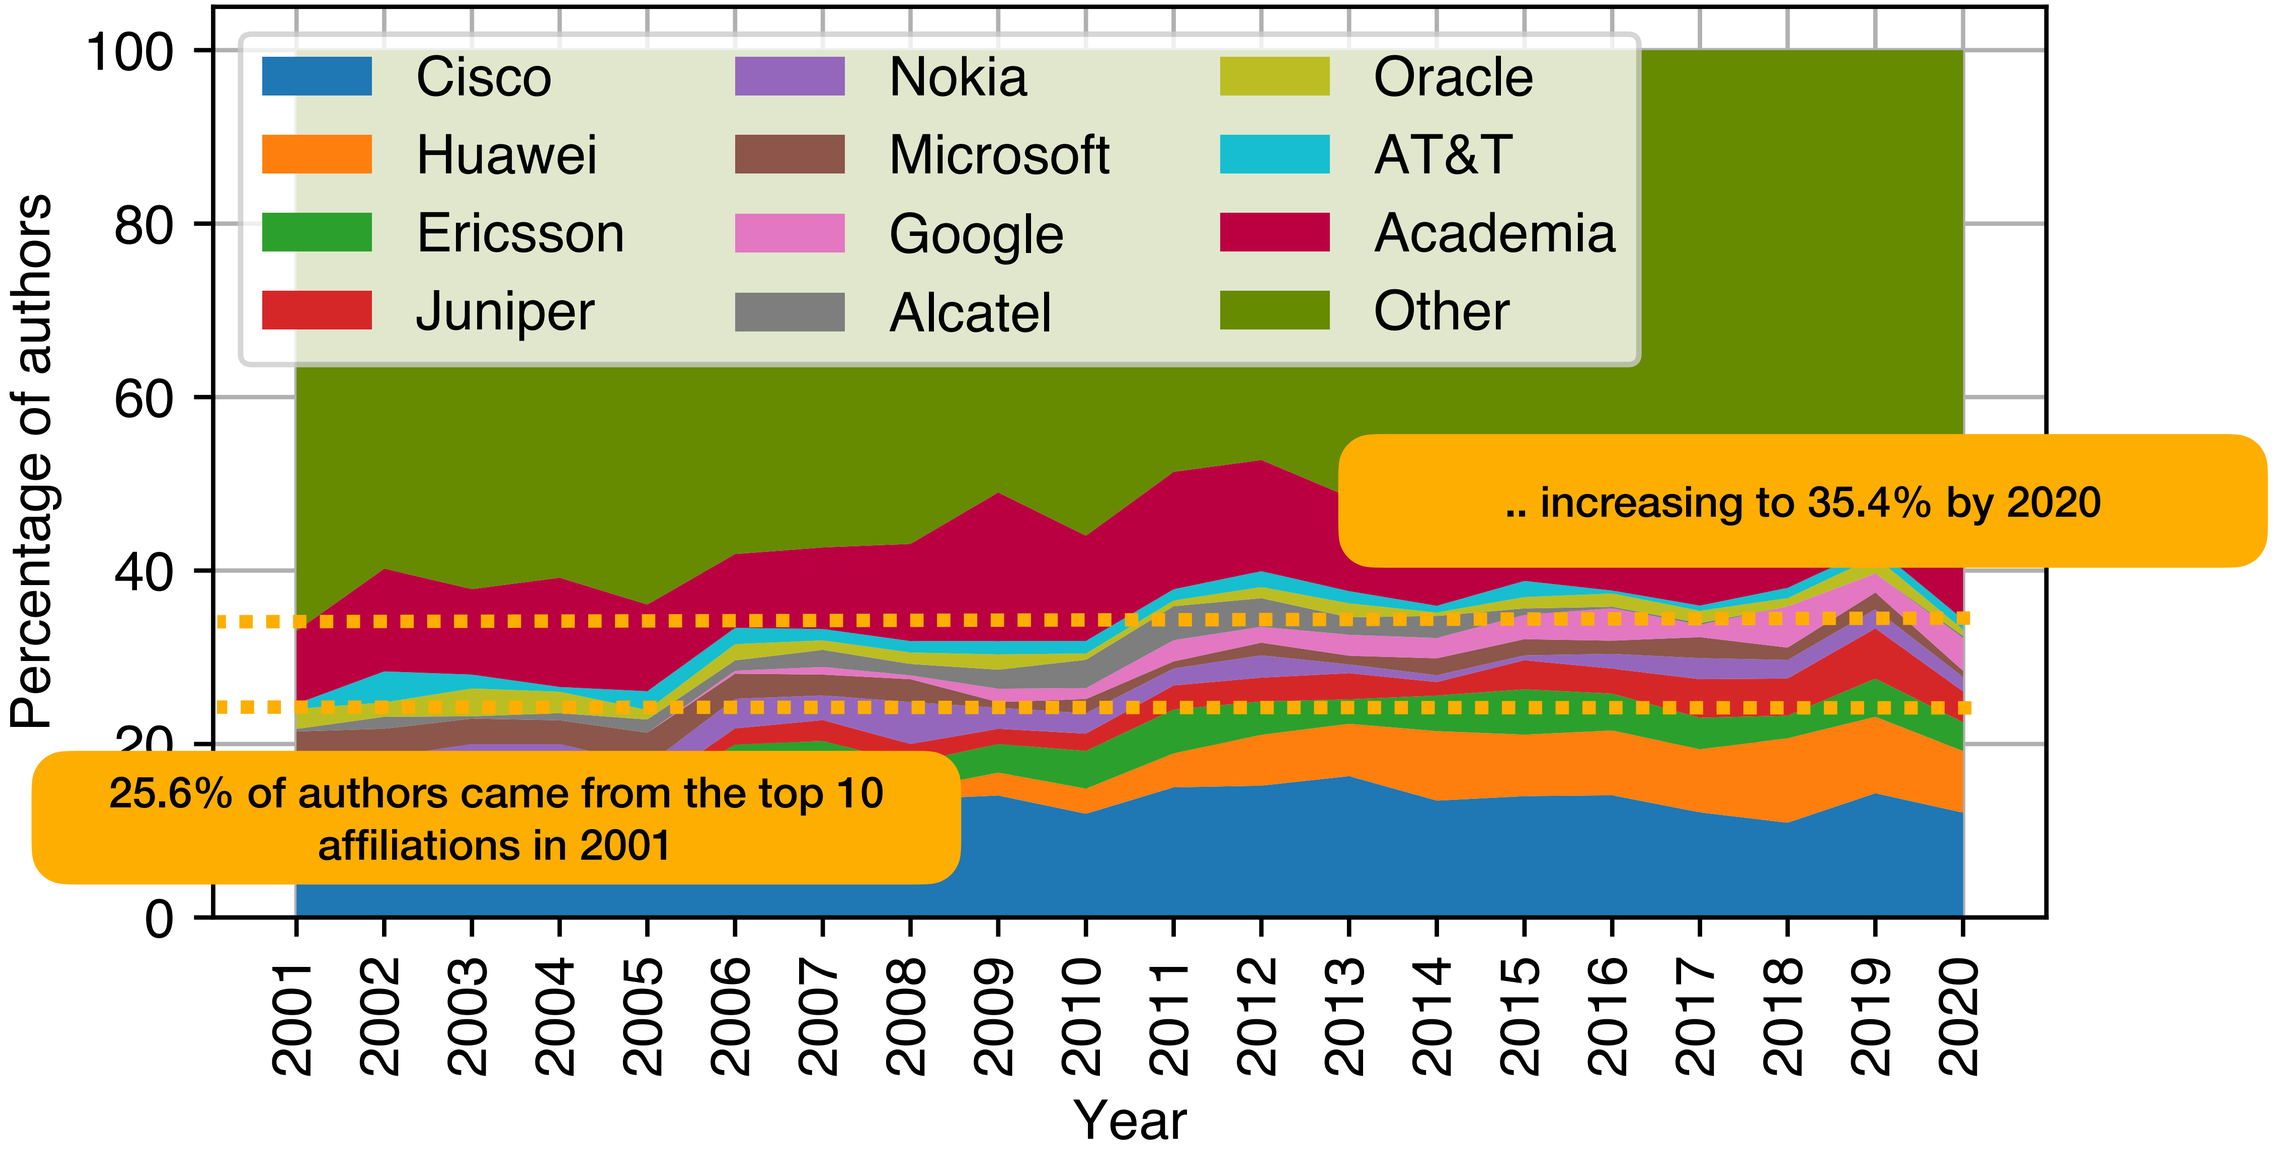
sodestream.github.io



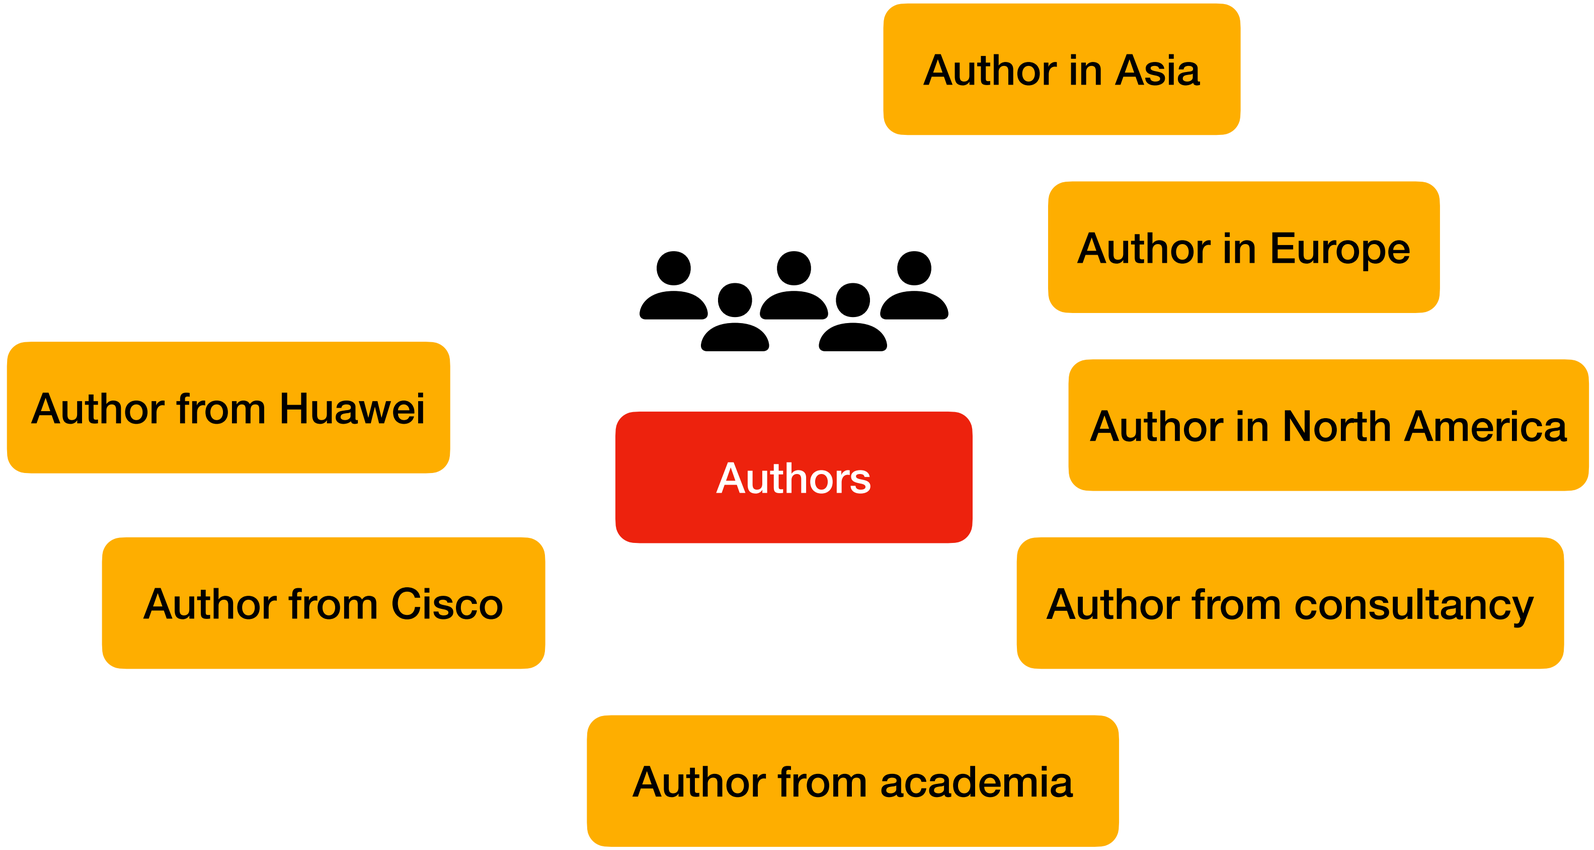
sodestream.github.io

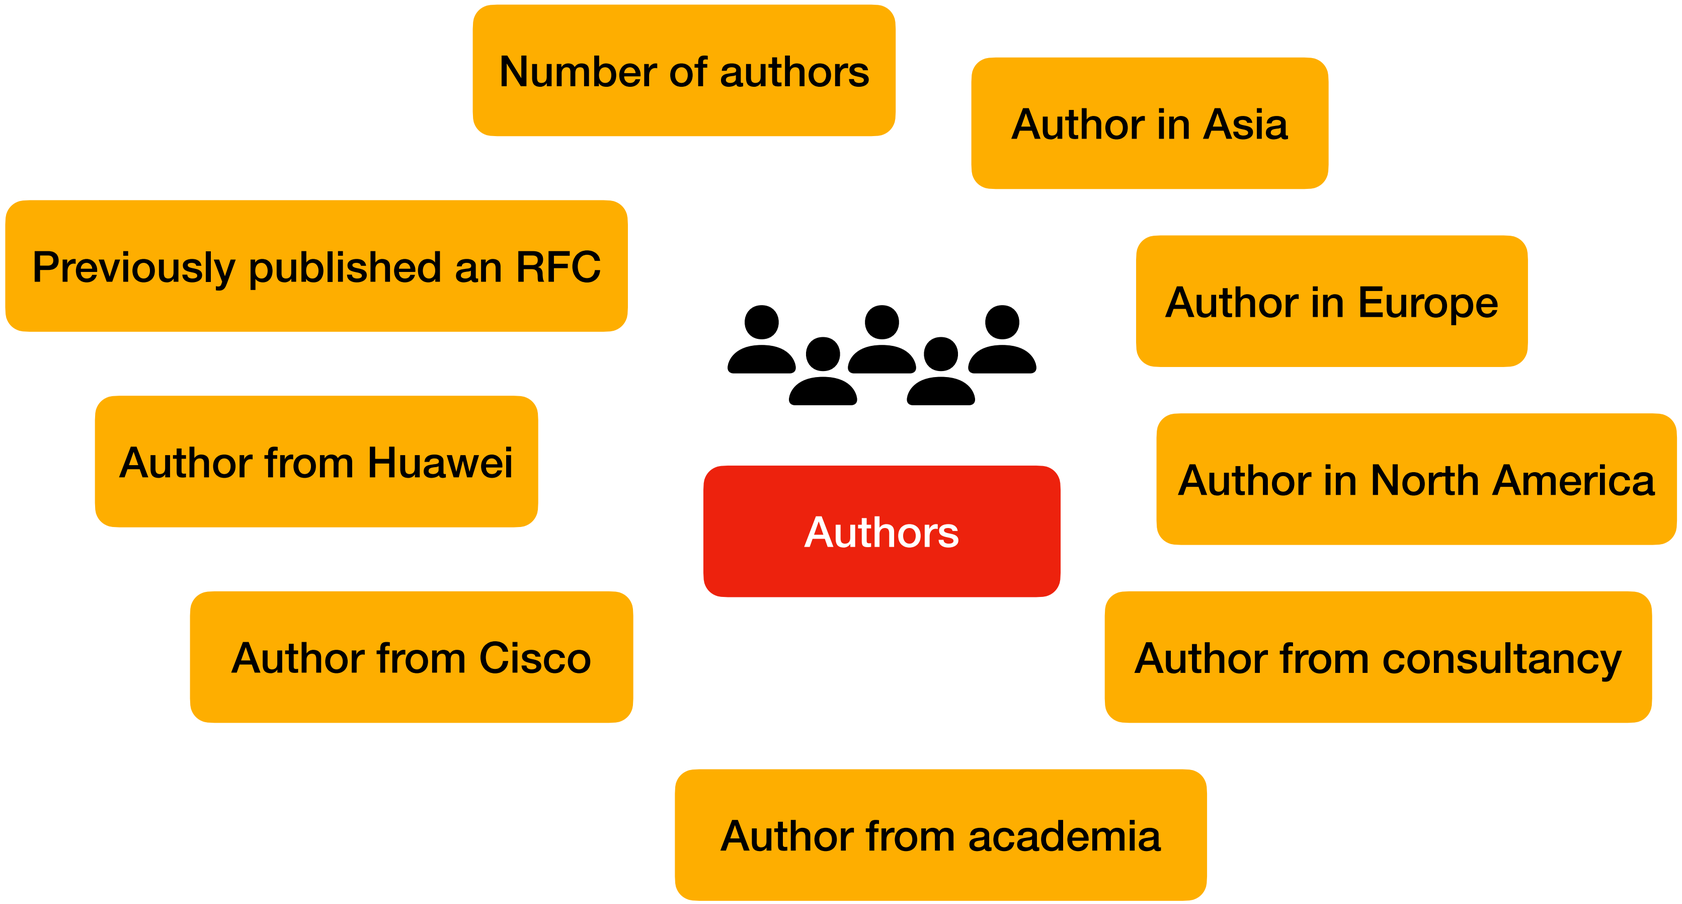


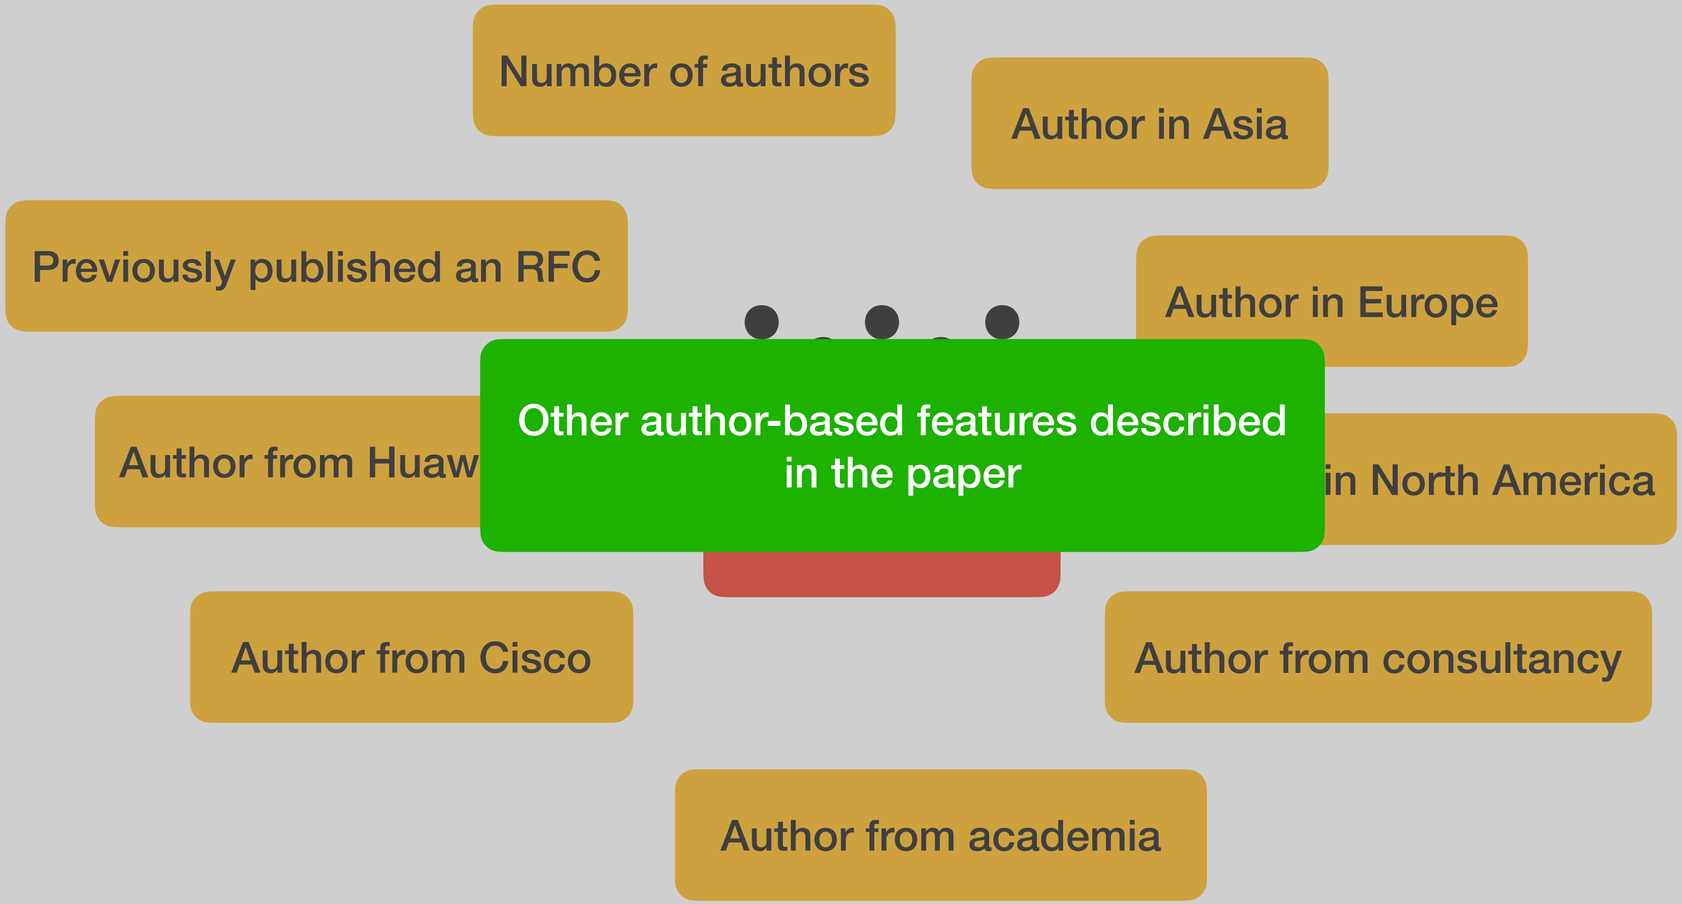
sodestream.github.io

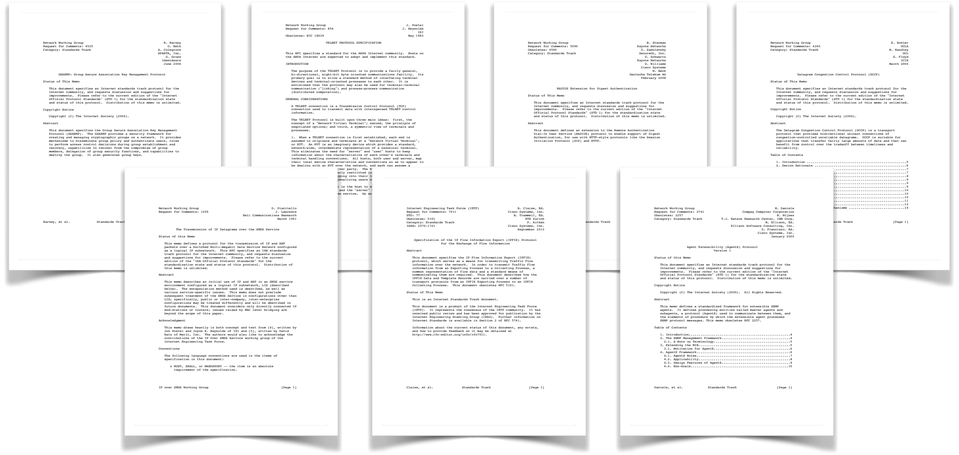


sodestream.github.io


Author in Asia
Author in Europe
Author from Huawei
Author in North America
Authors
Author from Cisco
Author from consultancy
Author from academia


Number of authors
Author in Asia
Previously published an RFC
Author in Europe
Author from Huawei
Author in North America
Authors
Author from Cisco
Author from consultancy
Author from academia


Number of authors
Author in Asia
Previously published an RFC
Author in Europe
Author from Huaw
Other author-based features described in the paper
in North America
Author from Cisco
Author from consultancy
Author from academia



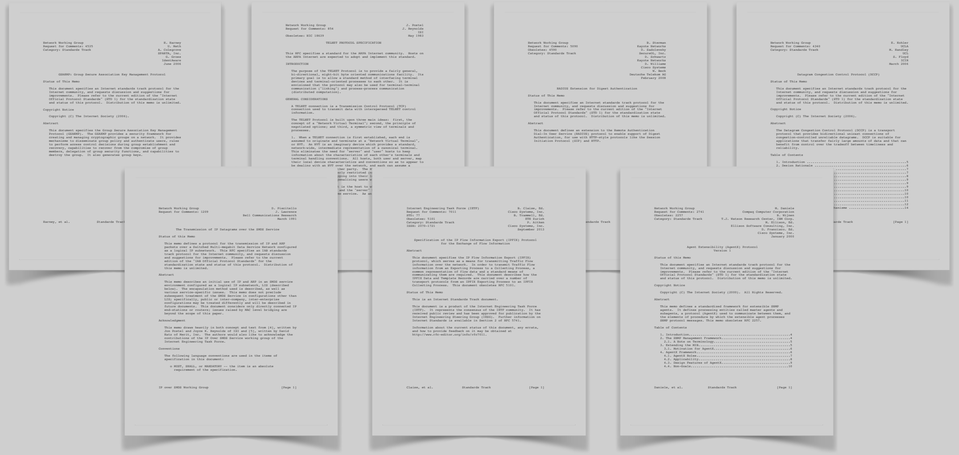
Documents

Authors

IETF®

Mailing lists

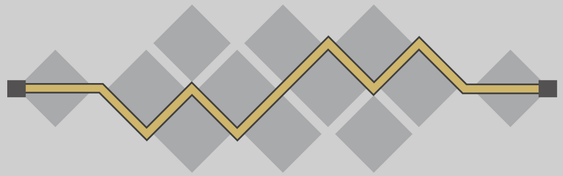
sodestream.github.io



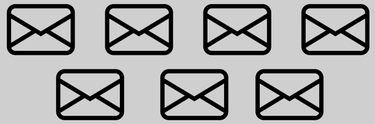
Documents



Authors

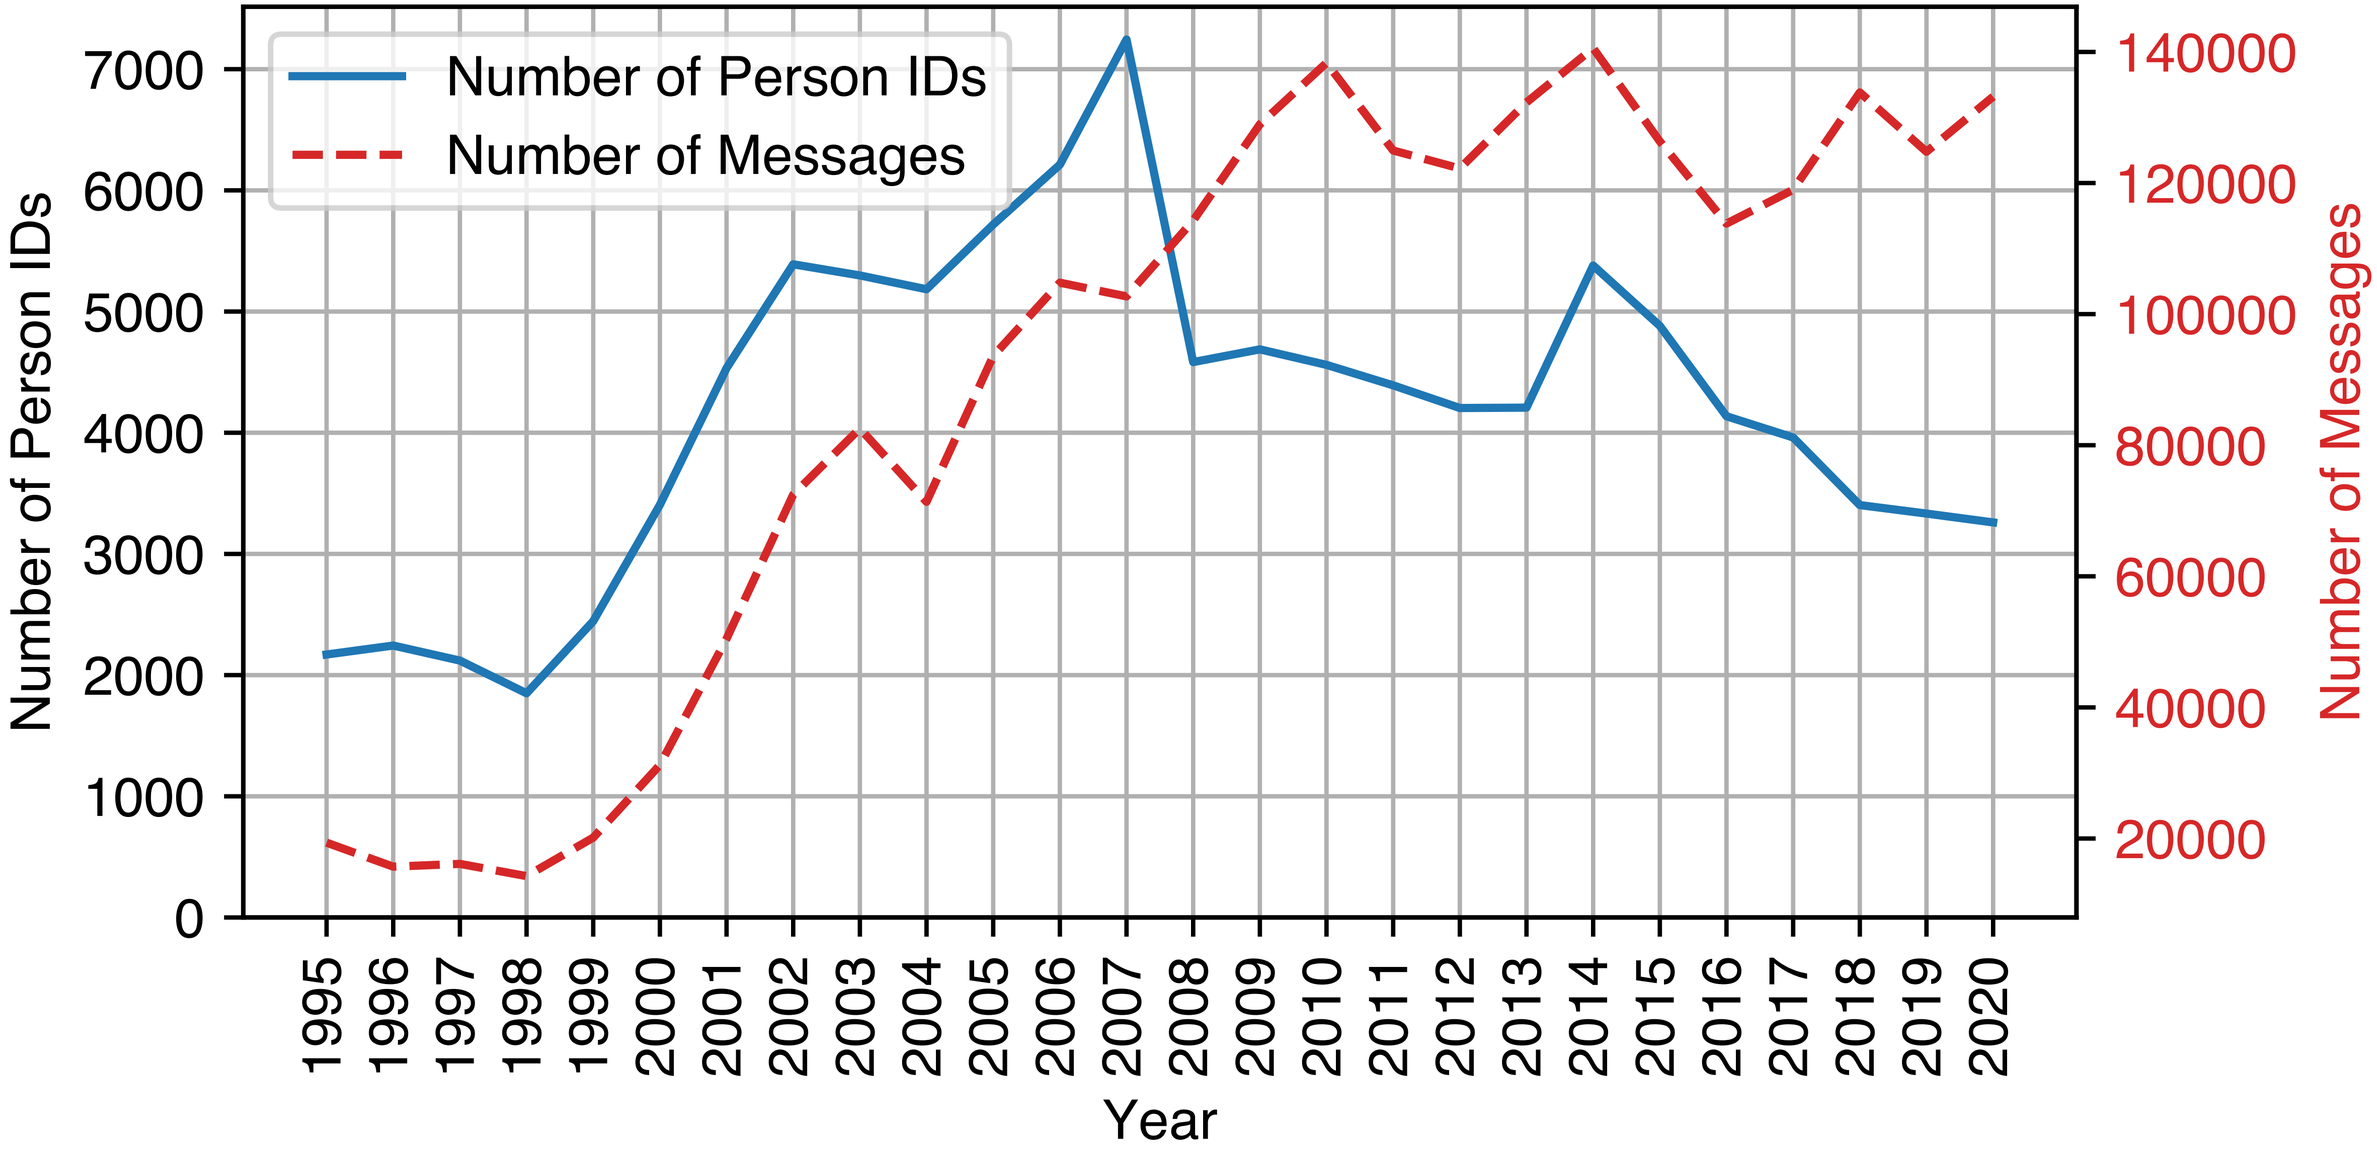


IETF®

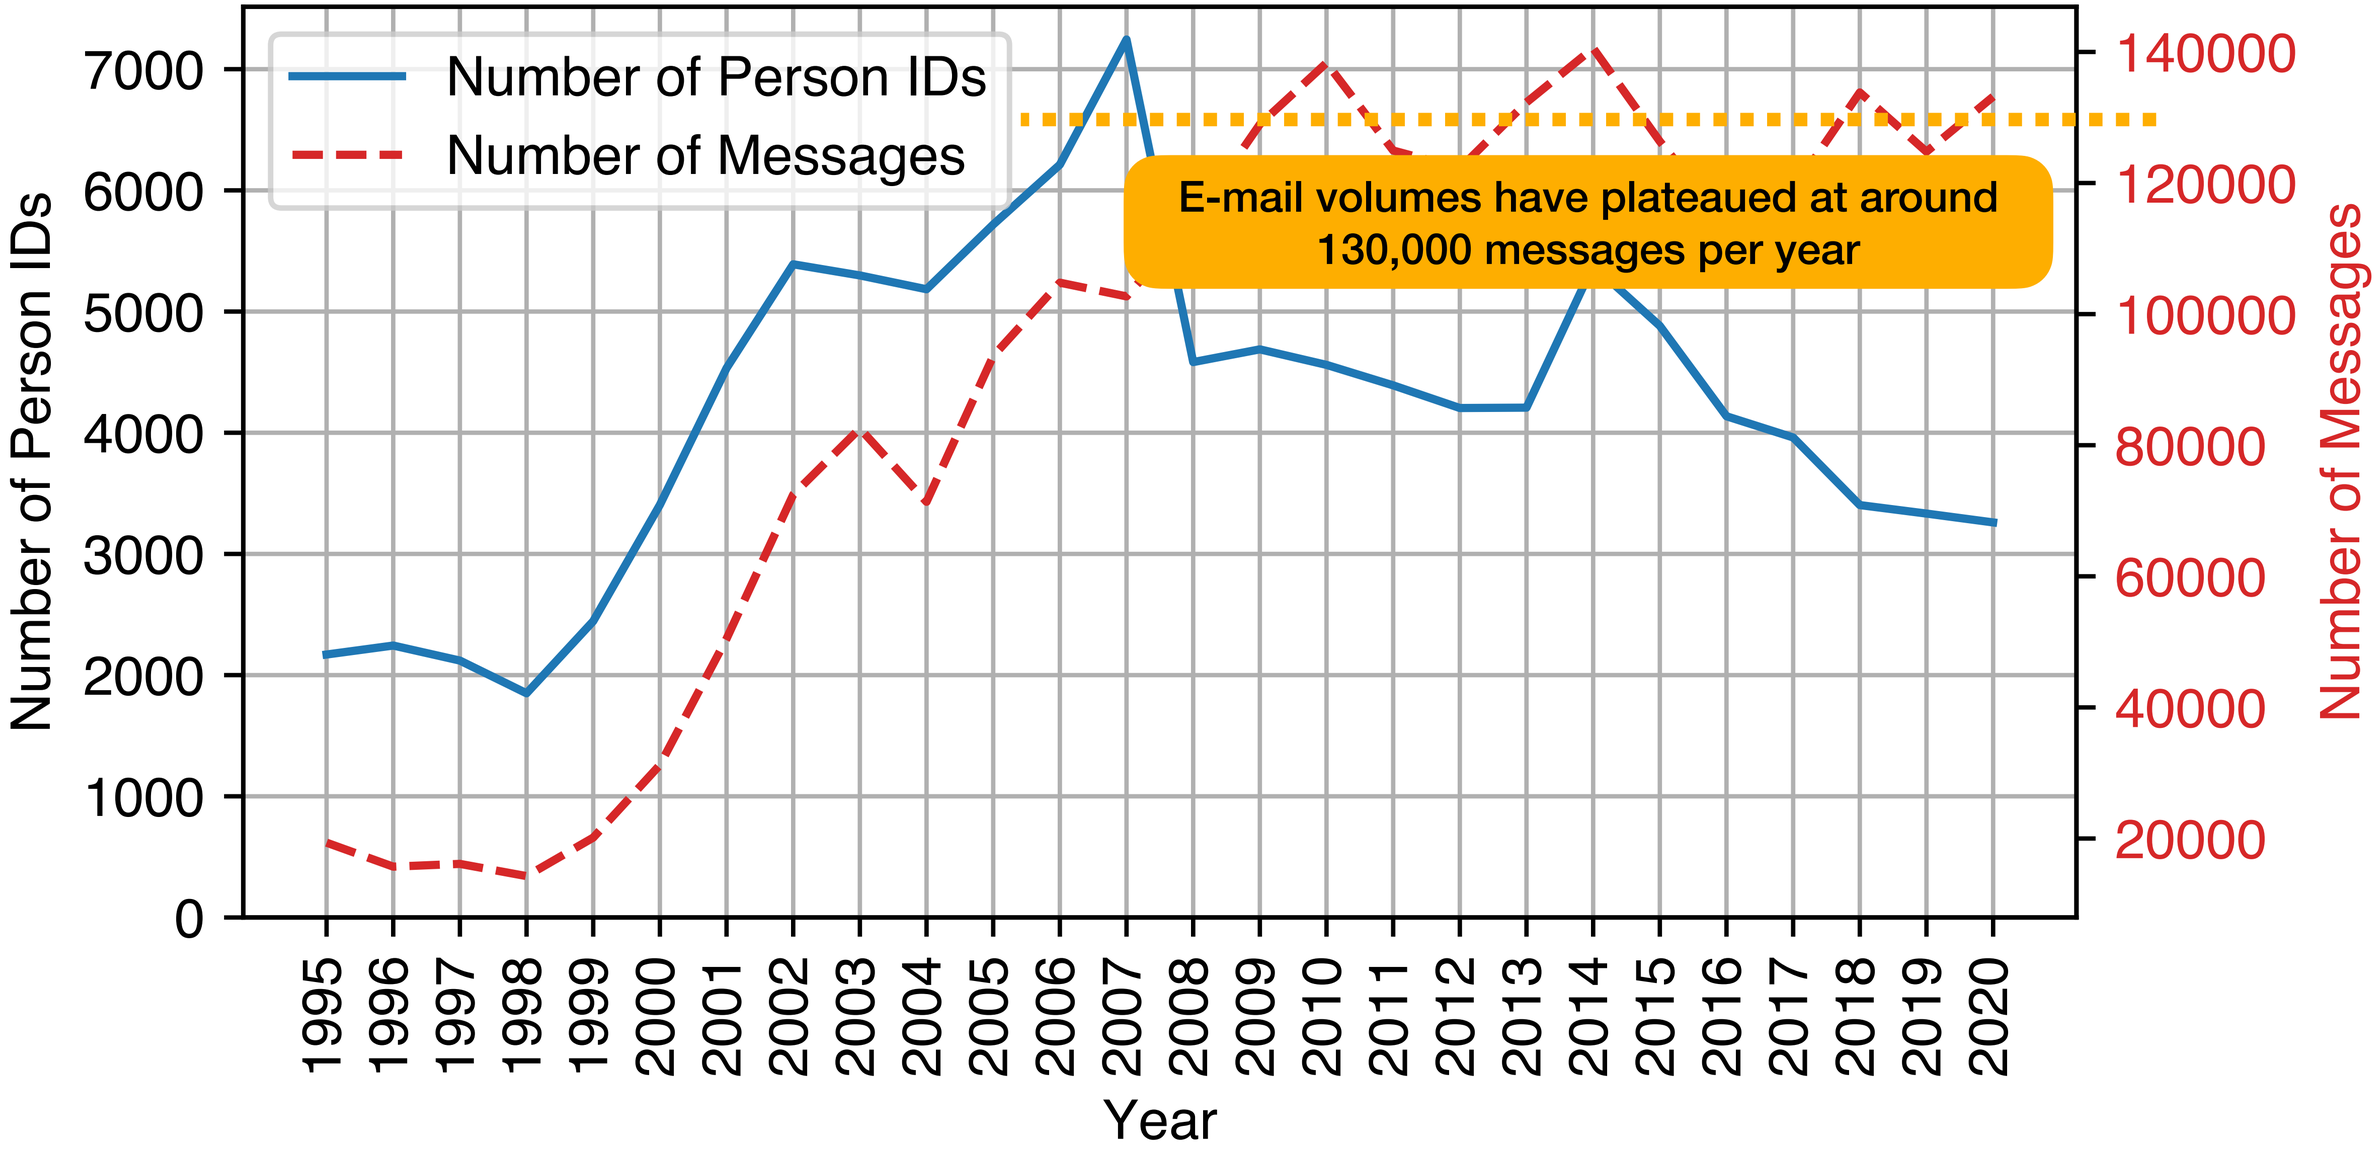


Mailing lists

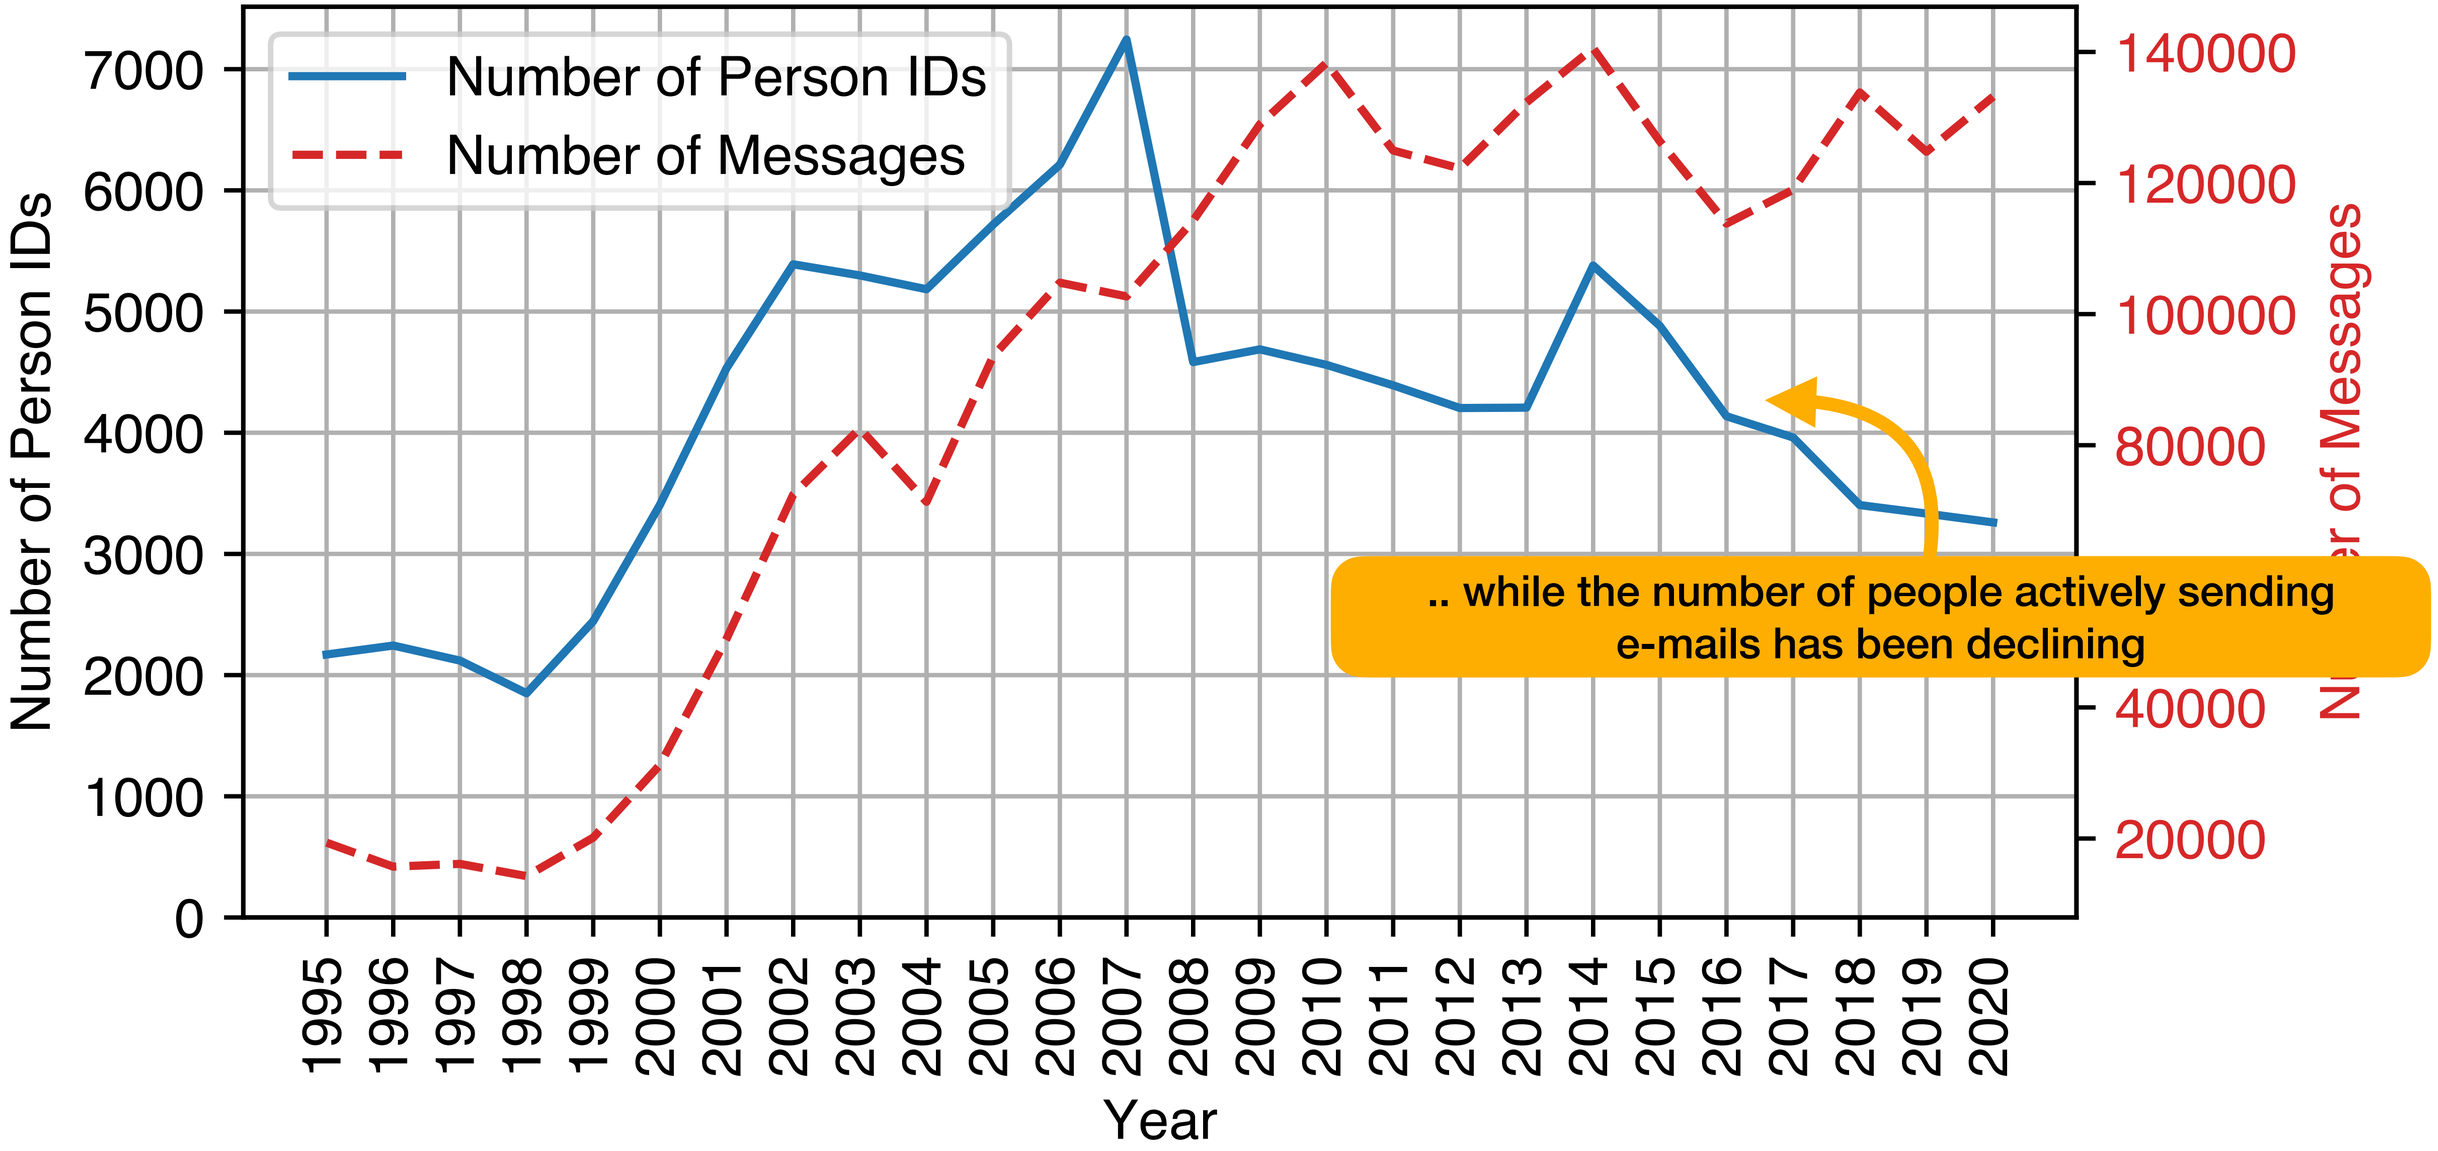
sodestream.github.io



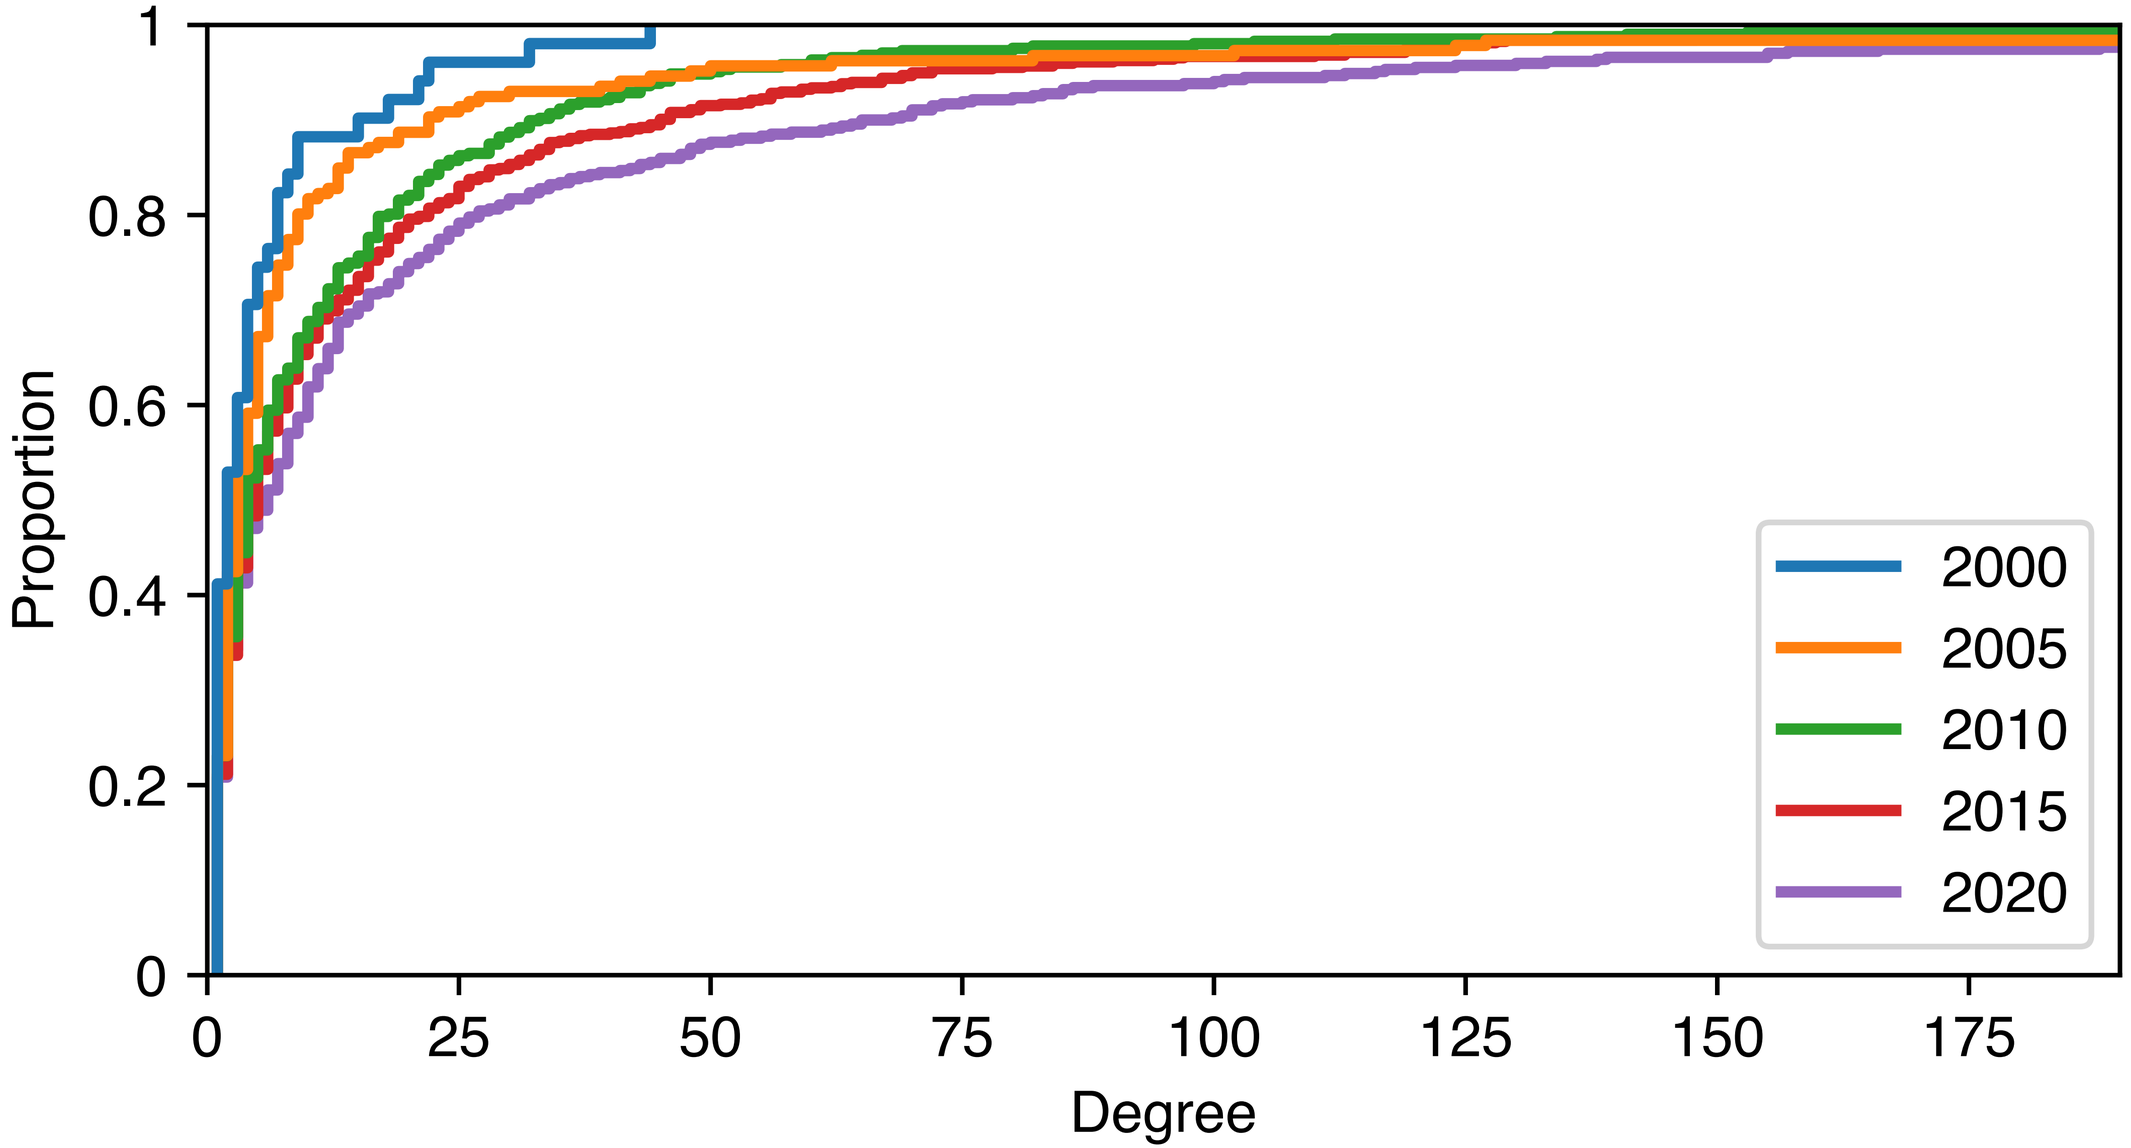
sodestream.github.io



sodestream.github.io

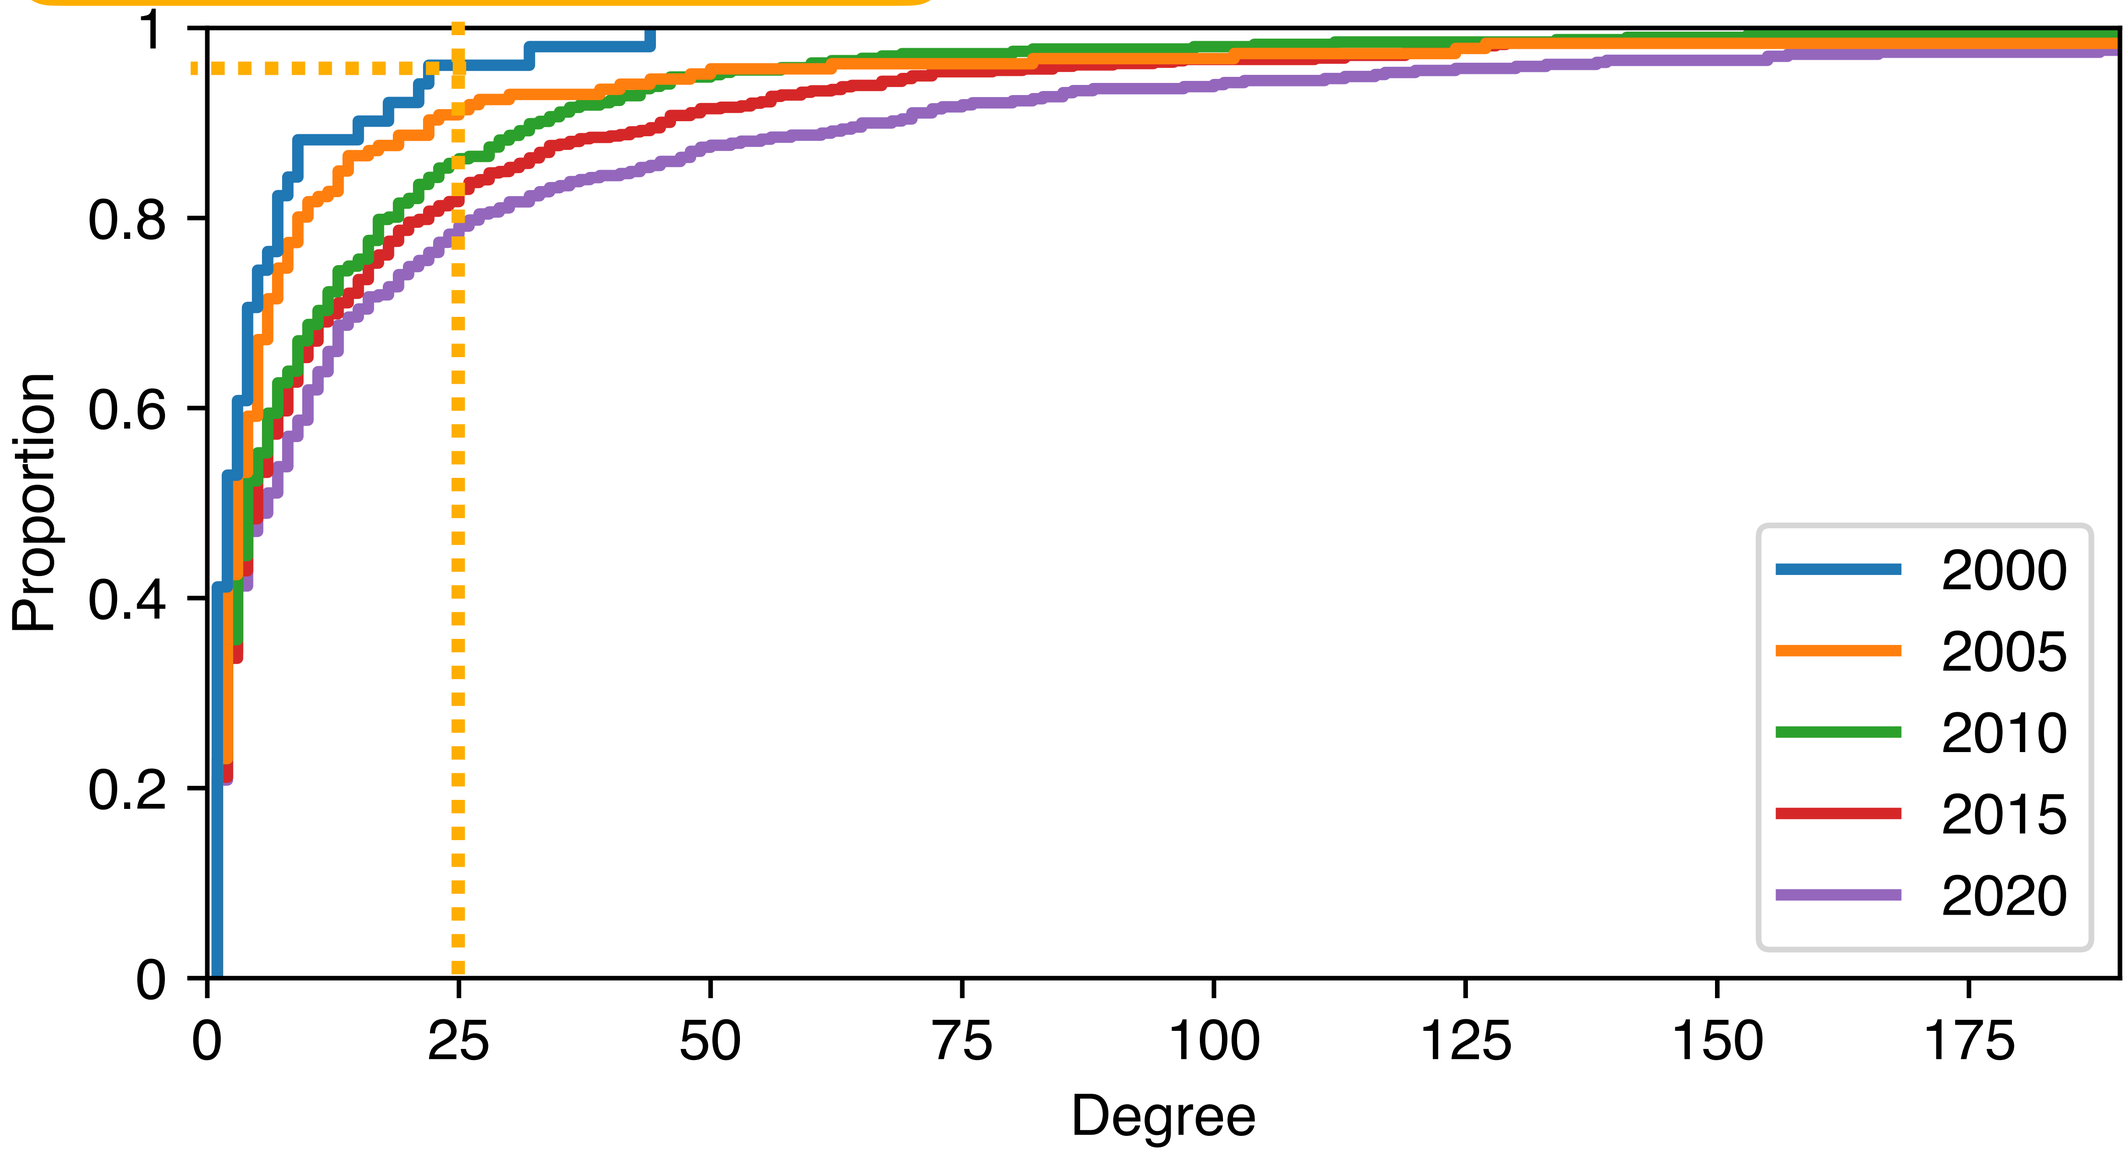


sodestream.github.io



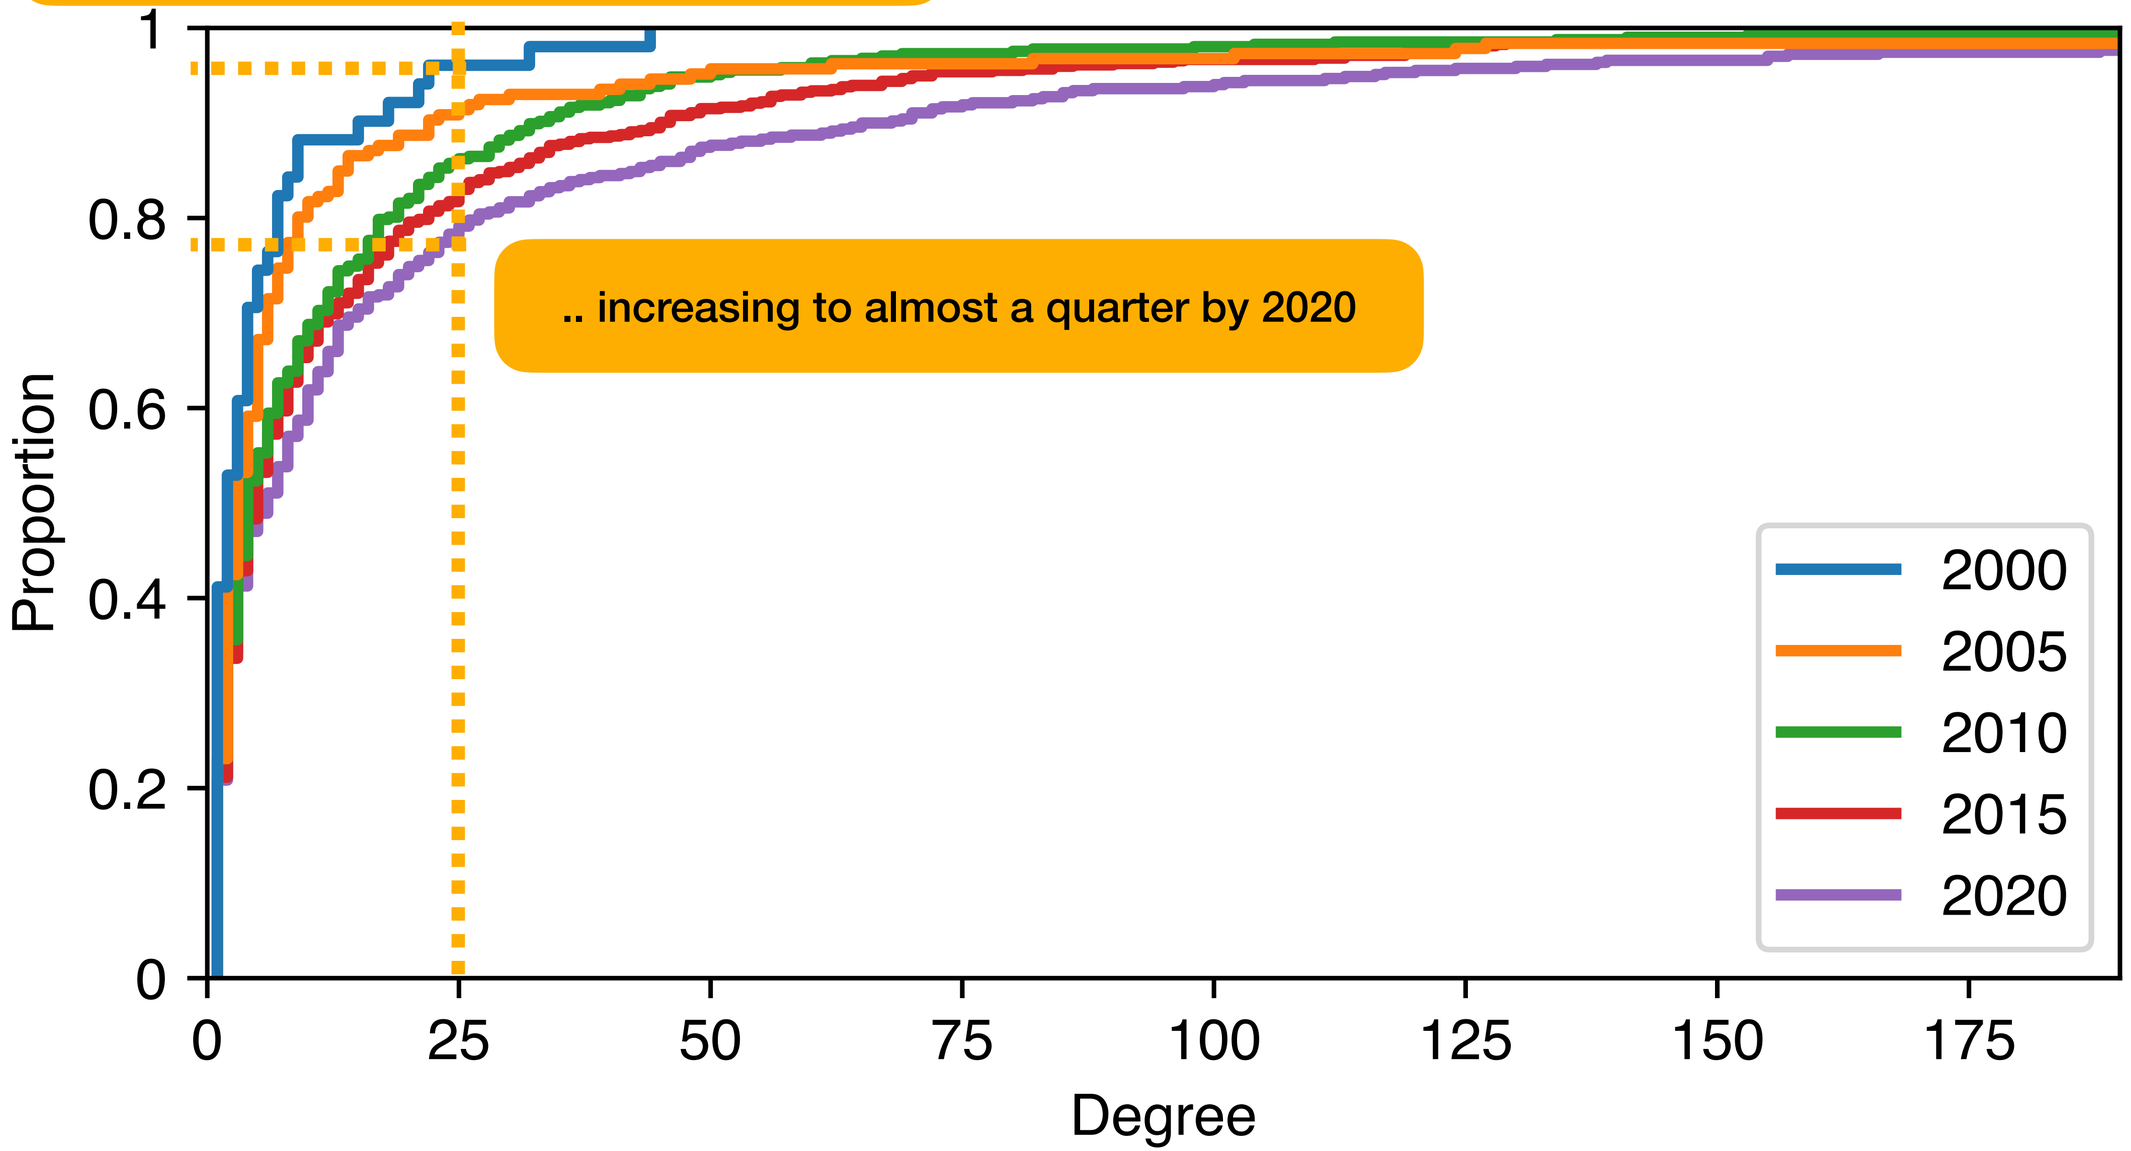
sodestream.github.io

Only 5.5% of authors interacted with more than 25 people in 2000



sodestream.github.io

Only 5.5% of authors interacted with more than 25 people in 2000

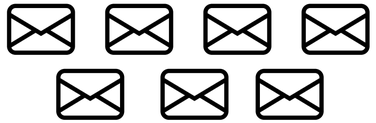


.. increasing to almost a quarter by 2020

sodestream.github.io

Number of e-mails mentioning drafts

Count of emails sent to junior author



Count of participants emailing junior author

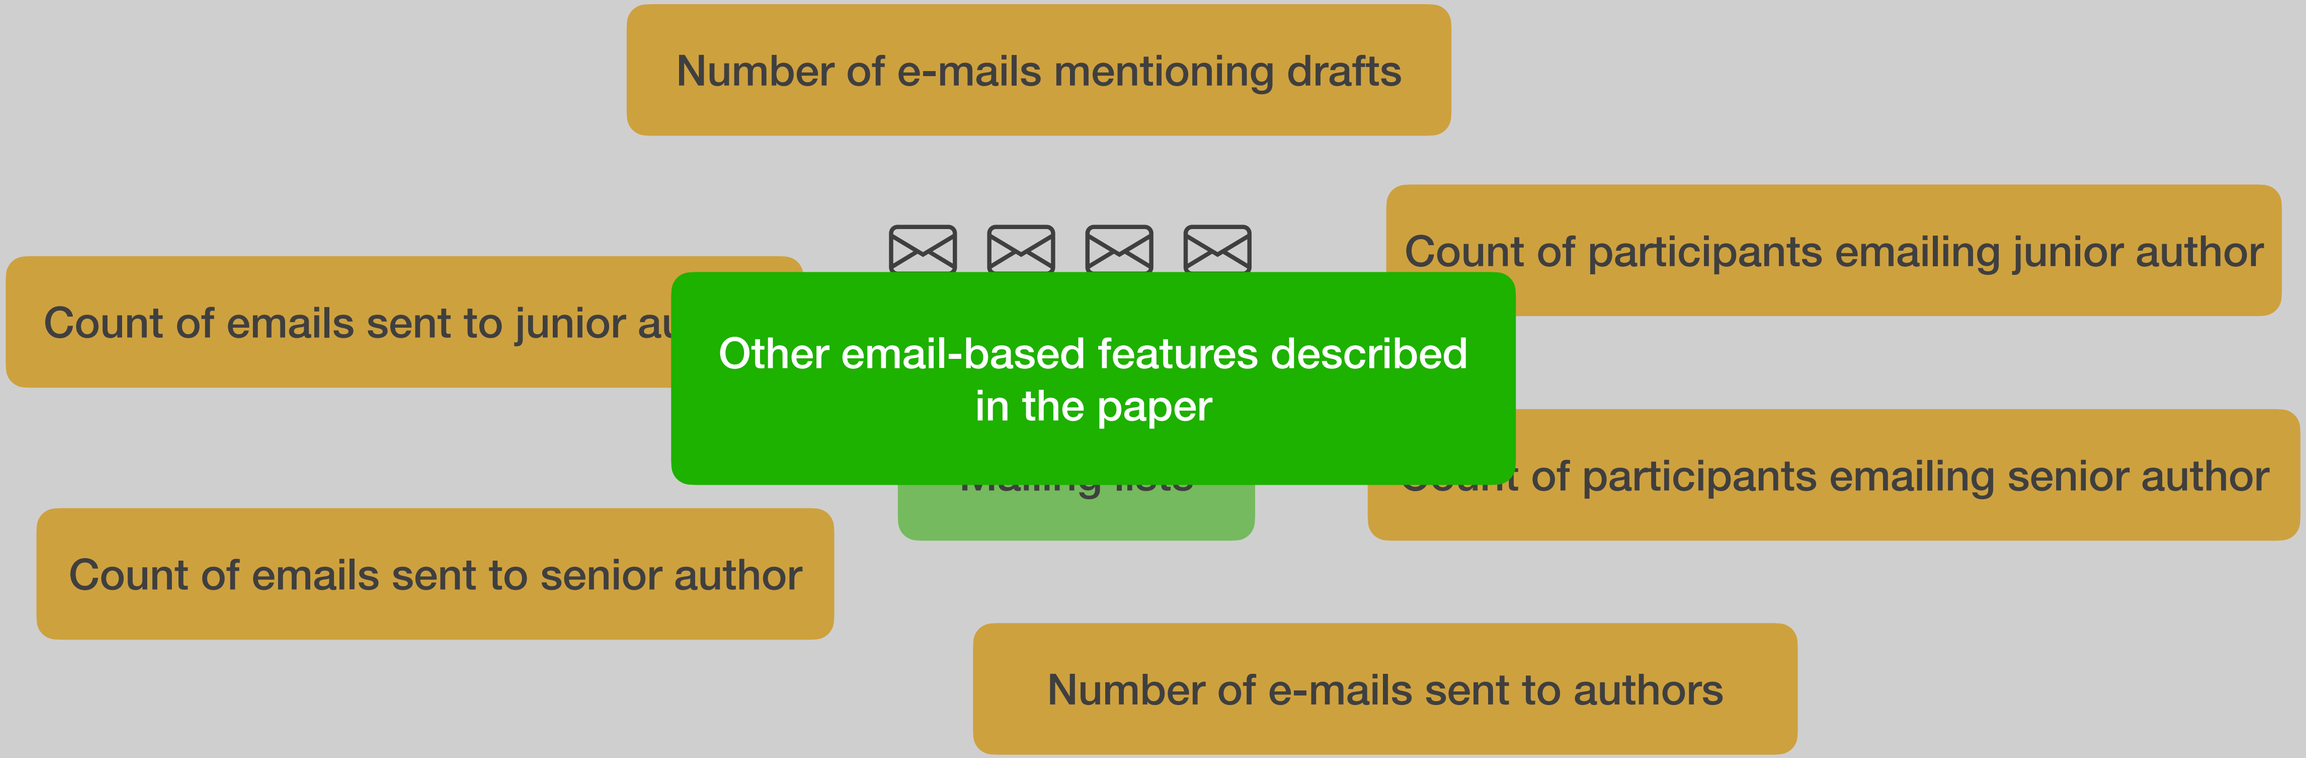
Mailing lists

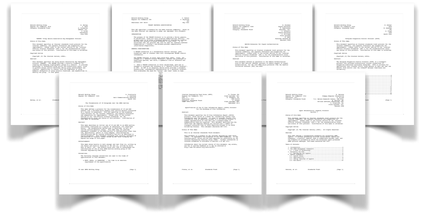
Count of participants emailing senior author

Count of emails sent to senior author

Number of e-mails sent to authors

sodestream.github.io

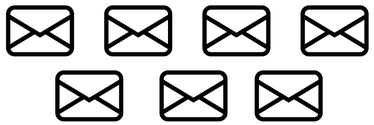

Number of e-mails mentioning drafts
Count of participants emailing junior author
Count of emails sent to junior au
Other email-based features described in the paper
Count of participants emailing senior author
Count of emails sent to senior author
Number of e-mails sent to authors



Documents

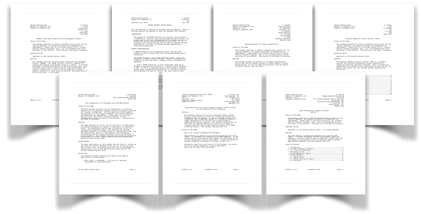


Authors



Mailing lists

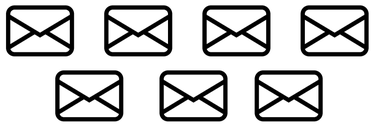
sodestream.github.io



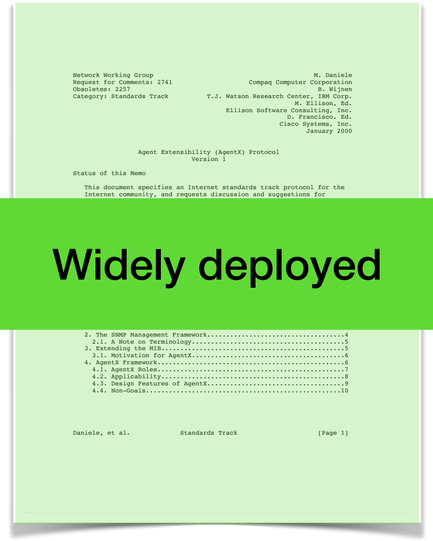
Documents



Authors



Mailing lists



Network Working Group M. Daniele
Request for Comments: 2741 Compaq Computer Corporation
Obsoletes: 2257 B. Wijnen
Category: Standards Track T.J. Watson Research Center, IBM Corp.
M. Ellison, Ed.
Ellison Software Consulting, Inc.
D. Francisco, Ed.
Cisco Systems, Inc.
January 2000

Agent Extensibility (AgentX) Protocol
Version 1

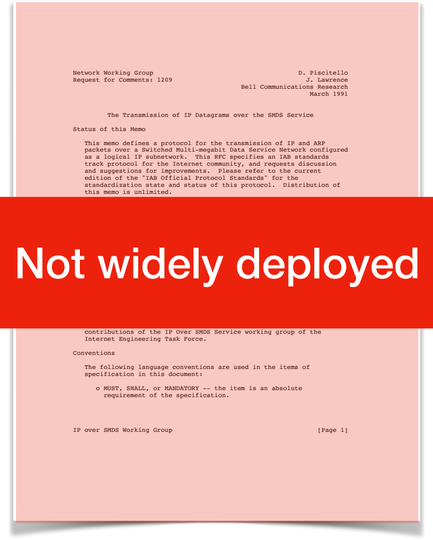
Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for

2. The SNMP Management Framework.....4
2.1. A Note on Terminology.....5
3. Extending the MIB.....5
3.1. Motivation for AgentX.....6
4. AgentX Framework.....6
4.1. AgentX Roles.....7
4.2. Applicability.....8
4.3. Design Features of AgentX.....9
4.4. Non-Goals.....10

Daniele, et al. Standards Track [Page 1]

Widely deployed



Network Working Group D. Piscitello
Request for Comments: 1209 J. Lawrence
Bell Communications Research
March 1991

The Transmission of IP Datagrams over the SMDS Service

Status of this Memo

This memo defines a protocol for the transmission of IP and ARP packets over a Switched Multi-megabit Data Service Network configured as a logical IP subnetwork. This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

contributions of the IP Over SMDS Service working group of the Internet Engineering Task Force.

Conventions

The following language conventions are used in the items of specification in this document:

o MUST, SHALL, or MANDATORY -- the item is an absolute requirement of the specification.

IP over SMDS Working Group [Page 1]

Not widely deployed

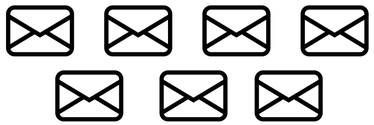
sodestream.github.io



Documents

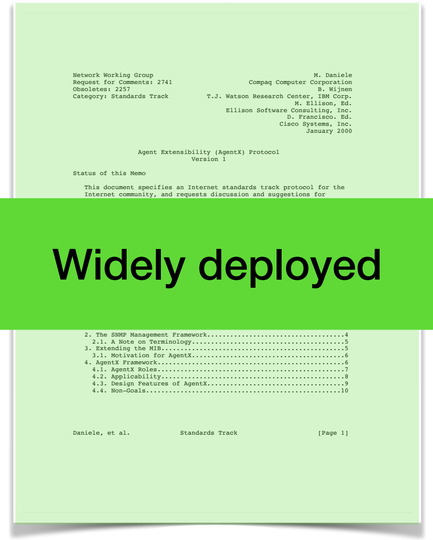


Authors

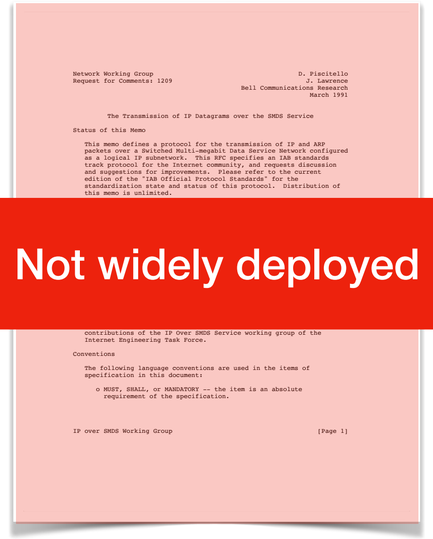


Mailing lists

Logistic regression model



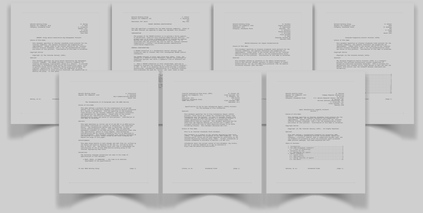
Widely deployed



Not widely deployed

sodestream.github.io

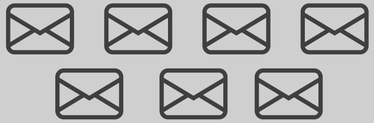
RFC deployment data from Nikkhah et al., “*A Statistical Exploration of Protocol Adoption*”. IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642



Documents

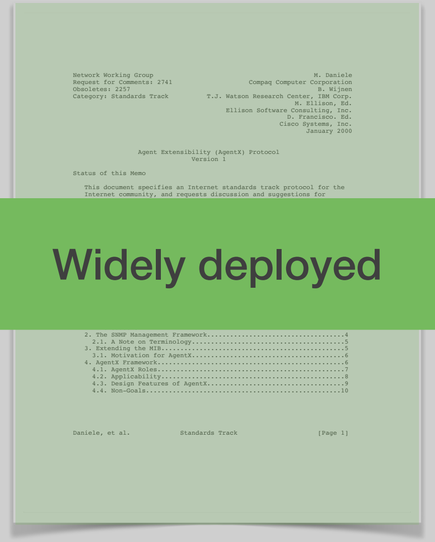


Authors

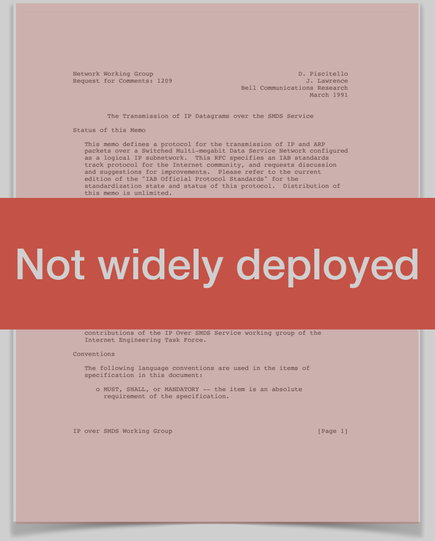


Mailing lists

More detailed discussion of the model in the paper



Widely deployed

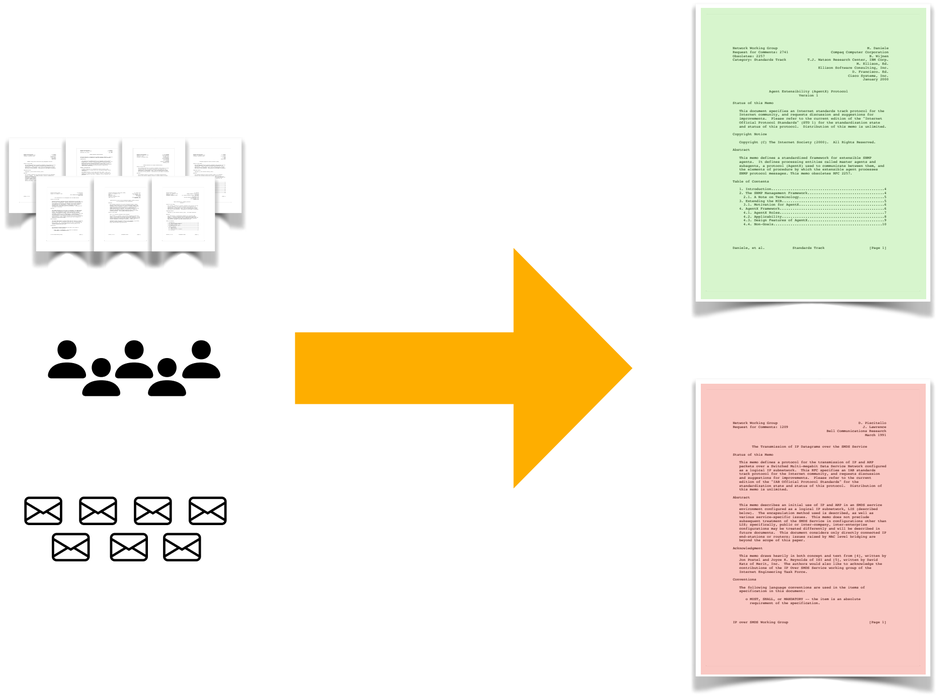


Not widely deployed

sodestream.github.io

Results

- Our additional features improved the performance of the model, vs. the baseline
- Features such as whether the RFC adds value, is scalable, builds on previous RFCs, is more widely cited by other RFCs, and includes more keywords are positively correlated with deployment
- We find that author demographic features, like geographic location and affiliation, aren’t significant predictors of deployment

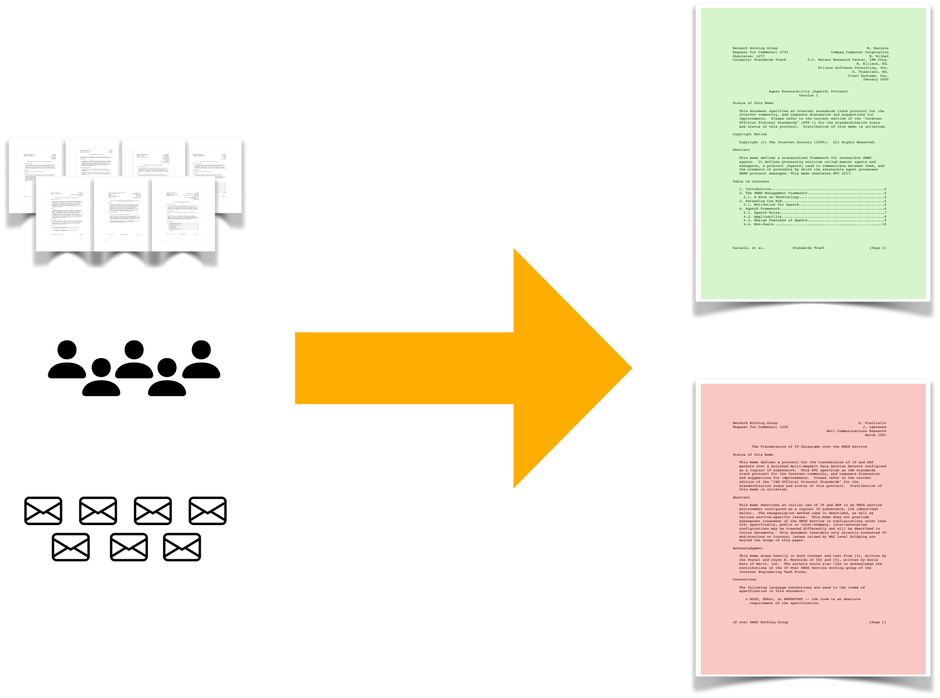


sodestream.github.io

Nikkhah et al., “*A Statistical Exploration of Protocol Adoption*”. IEEE/ACM Transactions on Networking. DOI: 10.1109/TNET.2017.2711642

Results

- Building on existing work
- Limited scope
- Setting normative requirements
- Meeting a need



sodestream.github.io

Summary

- RFC production rates are slowing, and RFCs are taking longer to produce
- The IETF's community is becoming more diverse, but issues remain
- The IETF itself does a good job of identifying successful protocols, and building on them
- Much more data left to analyse: GitHub issues, meeting materials, …



Characterising the IETF Through the Lens of RFC Deployment

Stephen McQuistin
sm@smcquistin.uk
University of Glasgow

Mladen Karan
m.karan@qmul.ac.uk
Queen Mary University of London

Prashant Khare
p.khare@qmul.ac.uk
Queen Mary University of London

Colin Perkins
csp@csperkins.org
University of Glasgow

Gareth Tyson
g.tyson@qmul.ac.uk
Queen Mary University of London

Matthew Purver
m.purver@qmul.ac.uk
Queen Mary University of London

Patrick Healey
p.healey@qmul.ac.uk
Queen Mary University of London

Waleed Iqbal
w.iqbal@qmul.ac.uk
Queen Mary University of London

Junaid Qadir
junaid.qadir@itu.edu.pk
Information Technology University

Ignacio Castro
i.castro@qmul.ac.uk
Queen Mary University of London

ABSTRACT

Protocol standards, defined by the Internet Engineering Task Force (IETF), are crucial to the successful operation of the Internet. This paper presents a large-scale empirical study of IETF activities, with a focus on understanding collaborative activities, and how these underpin the publication of standards documents (RFCs). Using a unique dataset of 2.4 million emails, 8,711 RFCs and 4,512 authors, we examine the shifts and trends within the standards development process, showing how protocol complexity and time to produce standards has increased. With these observations in mind, we develop statistical models to understand the factors that lead to successful uptake and deployment of protocols, deriving insights to improve the standardisation process.

CCS CONCEPTS

• **Social and professional topics** → **User characteristics**; • **Networks** → **Network protocol design**;

KEYWORDS

Protocol standardisation, IETF, Request for Comments

ACM Reference Format:
Stephen McQuistin, Mladen Karan, Prashant Khare, Colin Perkins, Gareth Tyson, Matthew Purver, Patrick Healey, Waleed Iqbal, Junaid Qadir, and Ignacio Castro. 2021. Characterising the IETF Through the Lens of RFC Deployment. In *ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3487552.3487821

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '21, November 2–4, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9129-0/21/11…$15.00
https://doi.org/10.1145/3487552.3487821

1 INTRODUCTION

Protocol standards are crucial to the successful operation of the Internet. A successful standard provides a basis for interoperability between systems developed by competing vendors, and supports the growth of an open ecosystem of products and services. Further, the process by which network protocol standards are developed, comprising multiple rounds of open feedback and review, has proven remarkably effective in designing high-quality and robust protocols, many of which see widespread deployment and use. Understanding the Internet standards development process, and how it produces successful protocols is, therefore, important if we are to understand the Internet and how it has evolved.

One of the main organisations that develops protocol standards is the Internet Engineering Task Force (IETF). The IETF was founded in 1986, following on from the US Government-funded effort that developed the early Internet. It has since grown to become a global community of network protocol designers, vendors, network operators, and researchers that develop and publish open network protocol standards and operational guidelines. The IETF publishes its standards, and other documents, in the RFC series (https://www.rfc-editor.org). This series comprises around 9,000 documents, authored over 50 years, and provides a rich history of the development of the Internet and its protocols [9].

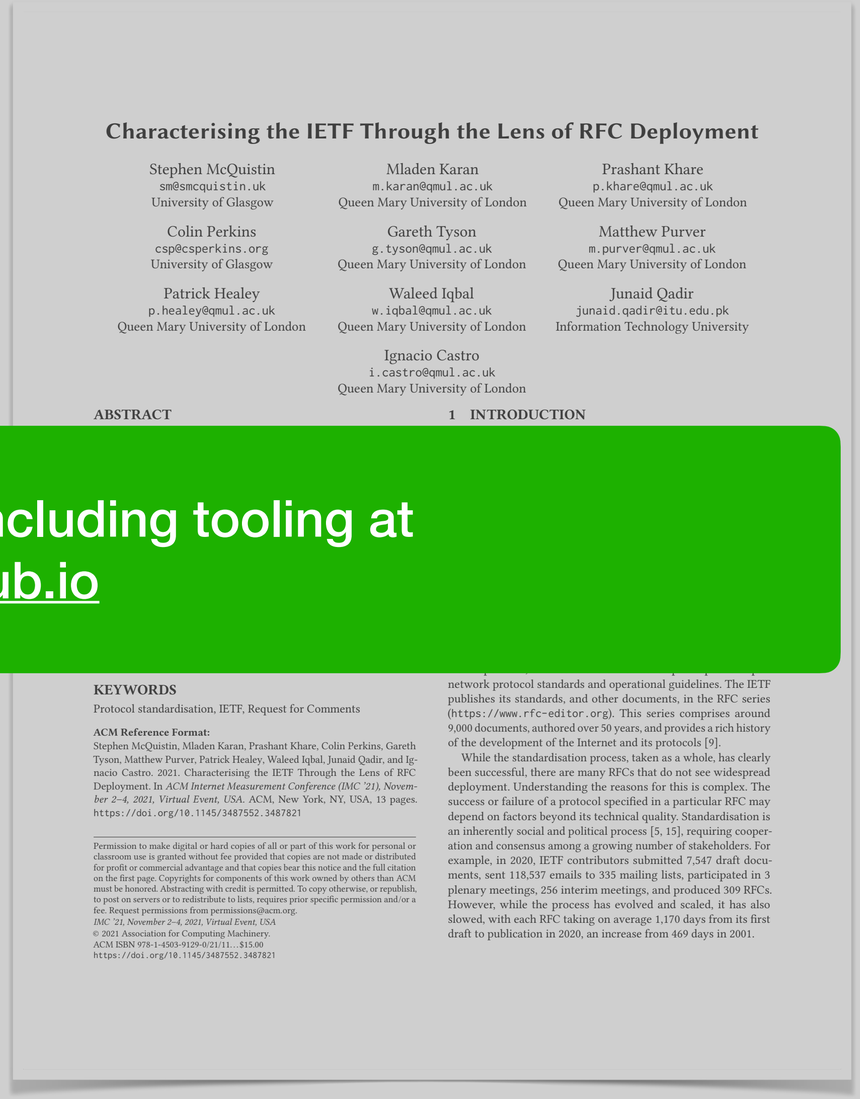
While the standardisation process, taken as a whole, has clearly been successful, there are many RFCs that do not see widespread deployment. Understanding the reasons for this is complex. The success or failure of a protocol specified in a particular RFC may depend on factors beyond its technical quality. Standardisation is an inherently social and political process [5, 15], requiring cooperation and consensus among a growing number of stakeholders. For example, in 2020, IETF contributors submitted 7,547 draft documents, sent 118,537 emails to 335 mailing lists, participated in 3 plenary meetings, 256 interim meetings, and produced 309 RFCs. However, while the process has evolved and scaled, it has also slowed, with each RFC taking on average 1,170 days from its first draft to publication in 2020, an increase from 469 days in 2001.

sodestream.github.io

Summary

- RFC production rates are slowing, and RFCs are taking longer to produce
- T
  d
- The IETF itself does a good job of identifying successful protocols, and building on them
- Much more data left to analyse: GitHub issues, meeting materials, …

More information about the project, including tooling at https://sodestream.github.io



Characterising the IETF Through the Lens of RFC Deployment

Stephen McQuistin
sm@smcquistin.uk
University of Glasgow

Mladen Karan
m.karan@qmul.ac.uk
Queen Mary University of London

Prashant Khare
p.khare@qmul.ac.uk
Queen Mary University of London

Colin Perkins
csp@csperkins.org
University of Glasgow

Gareth Tyson
g.tyson@qmul.ac.uk
Queen Mary University of London

Matthew Purver
m.purver@qmul.ac.uk
Queen Mary University of London

Patrick Healey
p.healey@qmul.ac.uk
Queen Mary University of London

Waleed Iqbal
w.iqbal@qmul.ac.uk
Queen Mary University of London

Junaid Qadir
junaid.qadir@itu.edu.pk
Information Technology University

Ignacio Castro
i.castro@qmul.ac.uk
Queen Mary University of London

ABSTRACT

1 INTRODUCTION

KEYWORDS

Protocol standardisation, IETF, Request for Comments

ACM Reference Format:
Stephen McQuistin, Mladen Karan, Prashant Khare, Colin Perkins, Gareth Tyson, Matthew Purver, Patrick Healey, Waleed Iqbal, Junaid Qadir, and Ignacio Castro. 2021. Characterising the IETF Through the Lens of RFC Deployment. In *ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3487552.3487821

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '21, November 2–4, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9129-0/21/11…$15.00
https://doi.org/10.1145/3487552.3487821

network protocol standards and operational guidelines. The IETF publishes its standards, and other documents, in the RFC series (https://www.rfc-editor.org). This series comprises around 9,000 documents, authored over 50 years, and provides a rich history of the development of the Internet and its protocols [9].

While the standardisation process, taken as a whole, has clearly been successful, there are many RFCs that do not see widespread deployment. Understanding the reasons for this is complex. The success or failure of a protocol specified in a particular RFC may depend on factors beyond its technical quality. Standardisation is an inherently social and political process [5, 15], requiring cooperation and consensus among a growing number of stakeholders. For example, in 2020, IETF contributors submitted 7,547 draft documents, sent 118,537 emails to 335 mailing lists, participated in 3 plenary meetings, 256 interim meetings, and produced 309 RFCs. However, while the process has evolved and scaled, it has also slowed, with each RFC taking on average 1,170 days from its first draft to publication in 2020, an increase from 469 days in 2001.

sodestream.github.io