



# Making the Internet: A Longitudinal Analysis of the IETF

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

Ignacio Castro Queen Mary University of London

**34th Multi-Service Networks Workshop, 8th July 2022**

 [sodestream.github.io](https://sodestream.github.io)



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  
T. Math  
A. Colegrove  
SPARTA, Inc.  
G. Gross  
IdentAware  
June 2006

GRAMP: 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 (GRAMP). The GRAMP 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  
Obsoletes: RFC 18639

J. Postel  
J. Reynolds  
ISI  
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  
Request for Comments: 5090  
Obsoletes: 4590  
Category: Standards Track

B. Stermann  
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  
4. Overview .....7  
5. Design .....8  
6. Design Rationale .....9  
7. Conventions and Terminology .....9  
8. Overview .....10  
9. Design .....10  
10. Design Rationale .....11  
11. Conventions and Terminology .....11  
12. Overview .....12  
13. Design .....12  
14. Design Rationale .....12

Harney, et al. Standards Track

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:  
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)  
Request for Comments: 7011  
SPD: 77  
Obsoletes: 5101  
Category: Standards Track  
ISBN: 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>.

Claise, et al. Standards Track [Page 1]


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


M. Daniele  
Compag 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

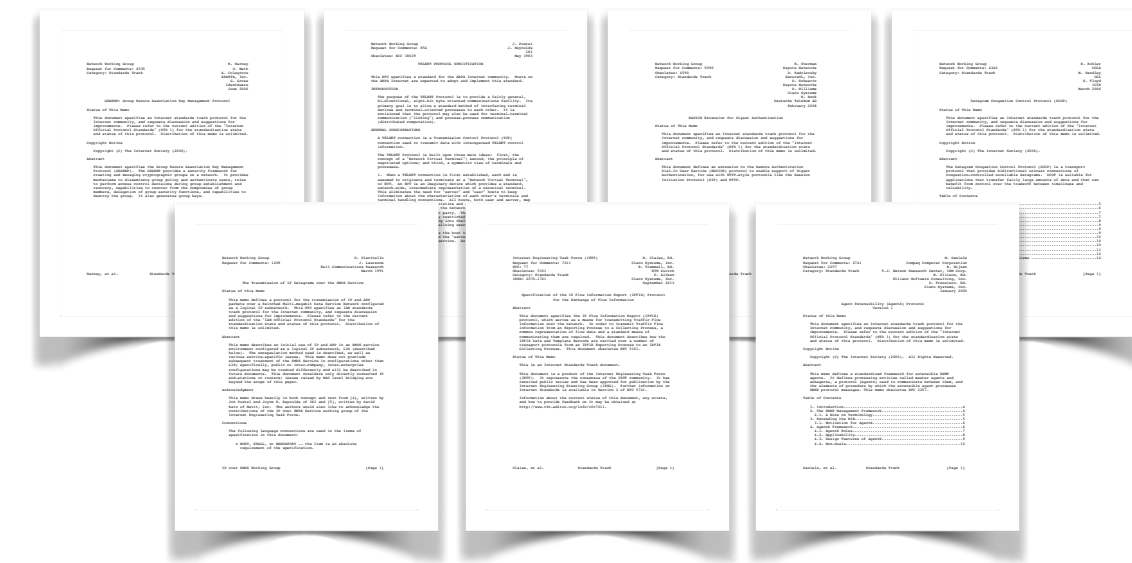
Status of this Memo  
This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of 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 of AgentX .....8  
4.3. Design Features of AgentX .....9  
4.4. Non-Goals .....10


Daniele, et al. Standards Track [Page 1]






RFCs


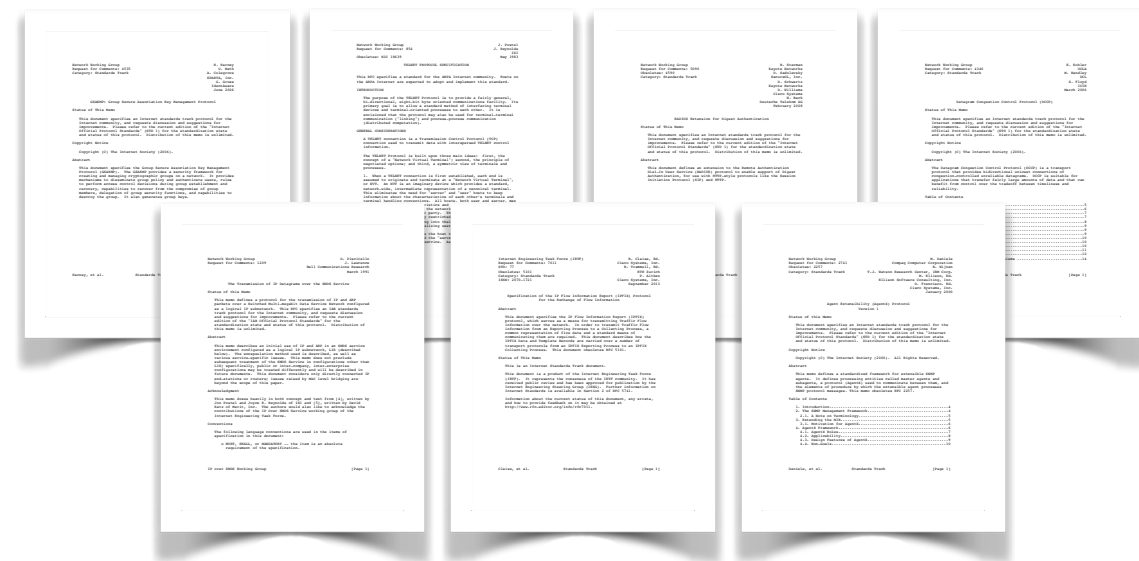




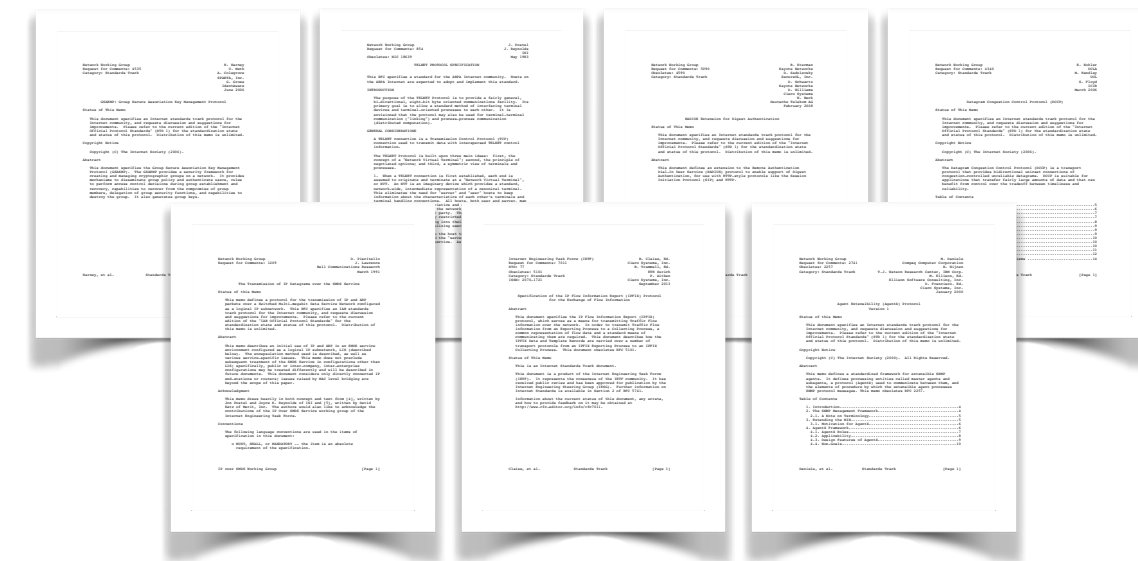
Authors

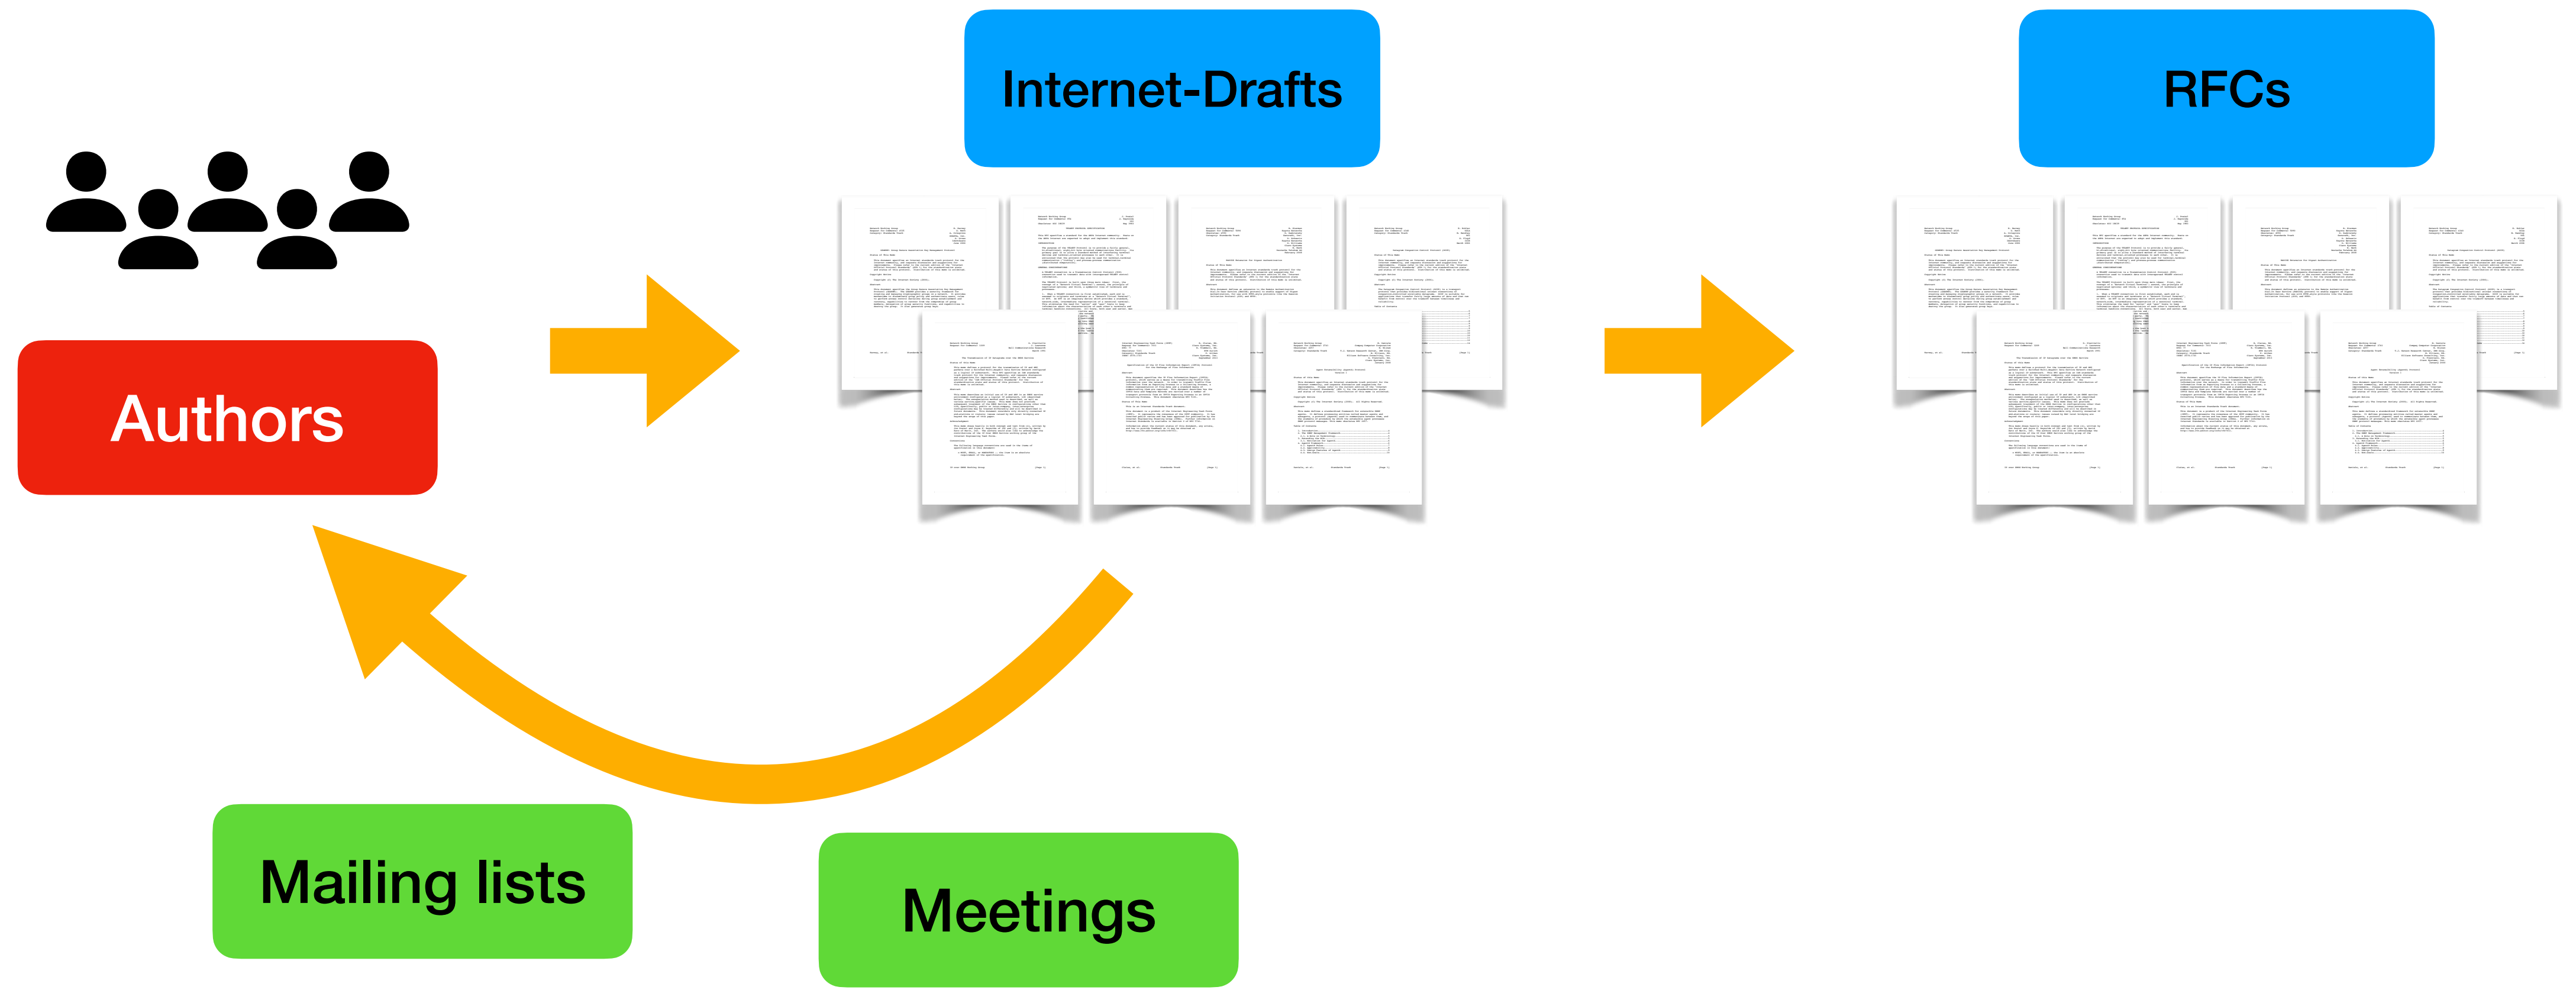


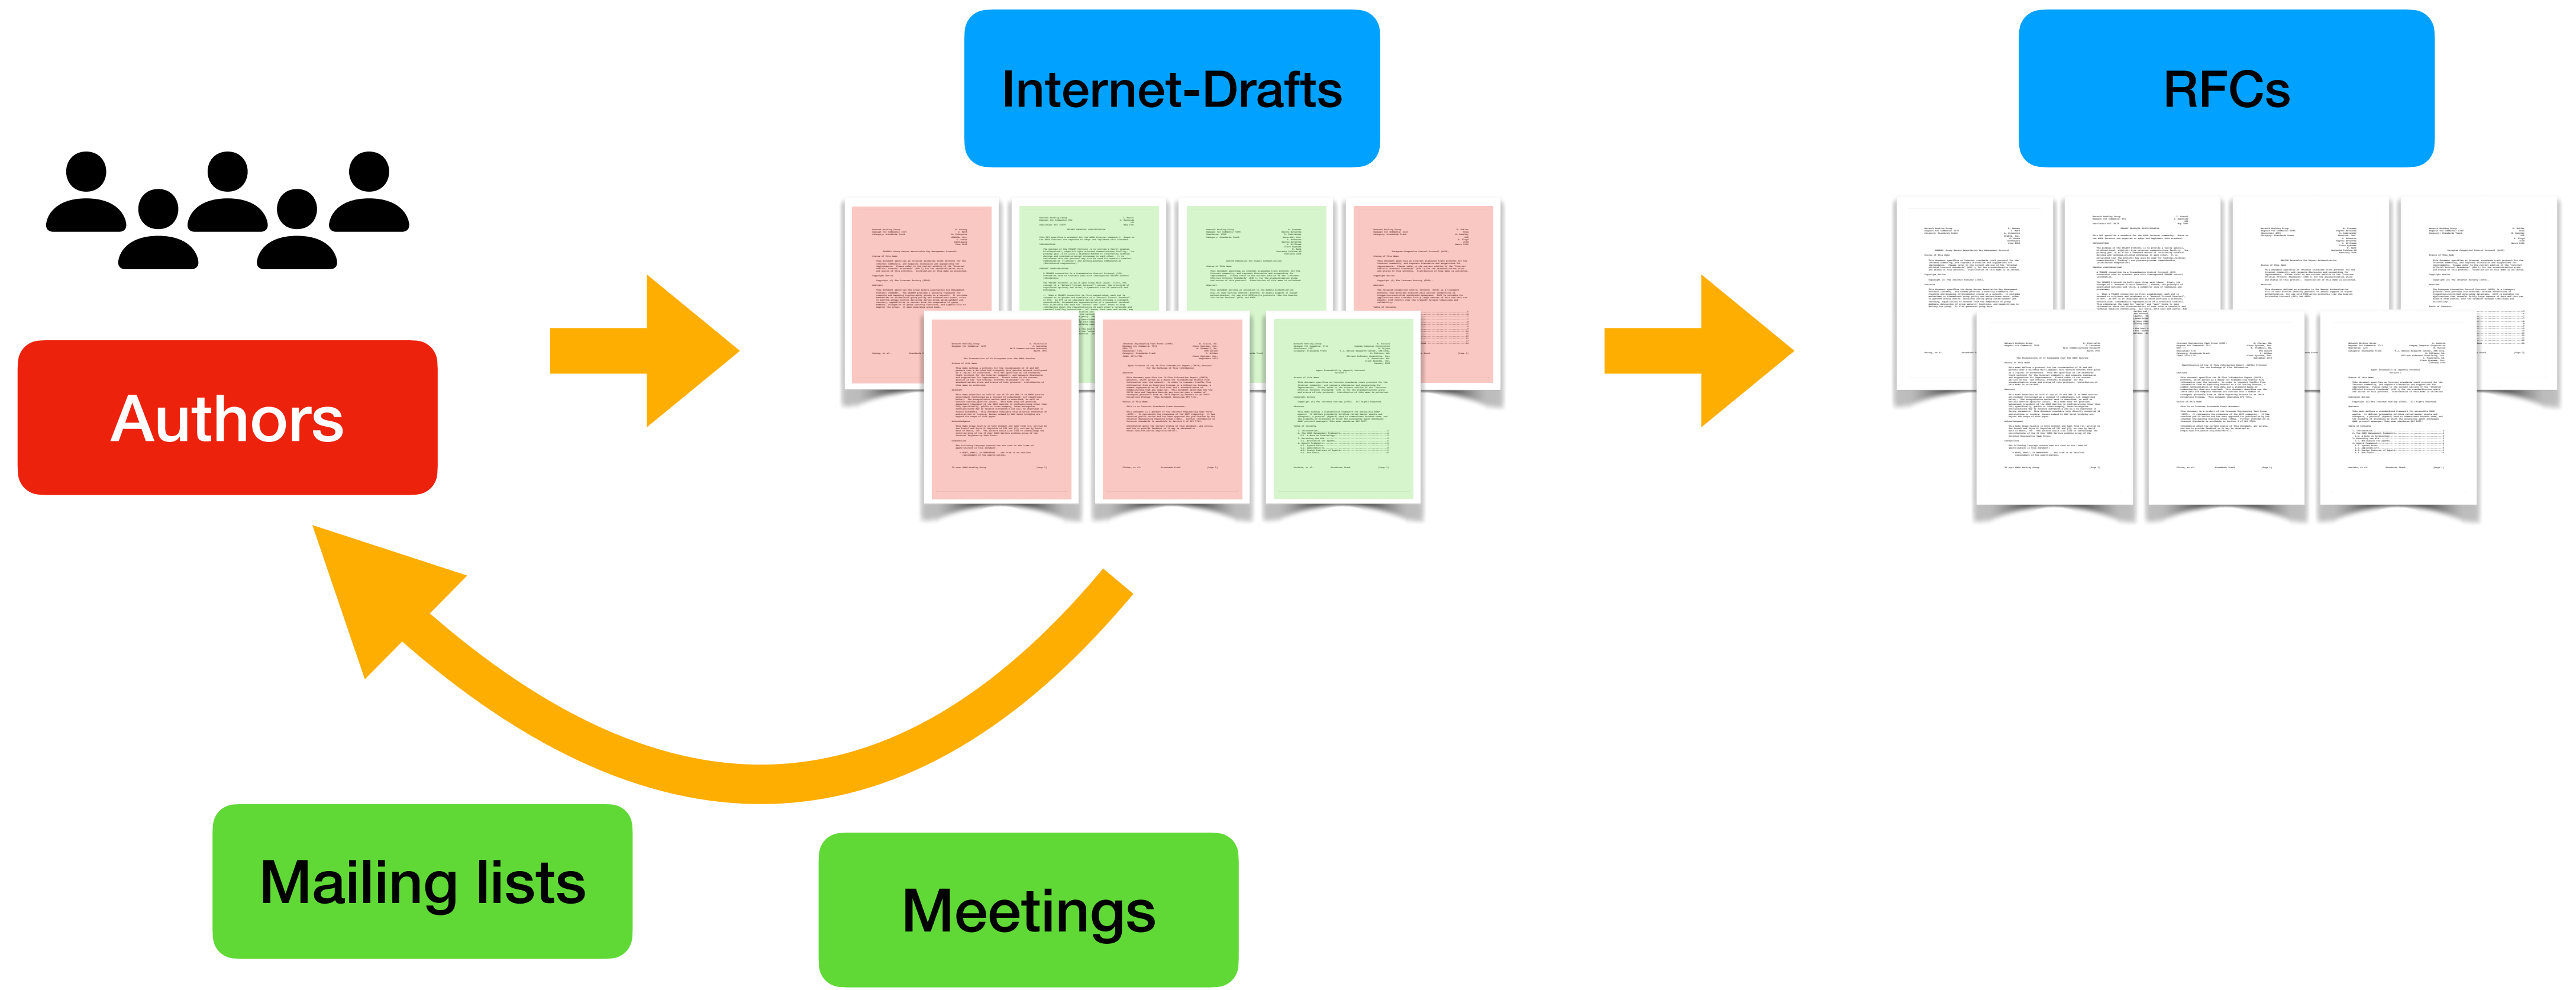
Internet-Drafts

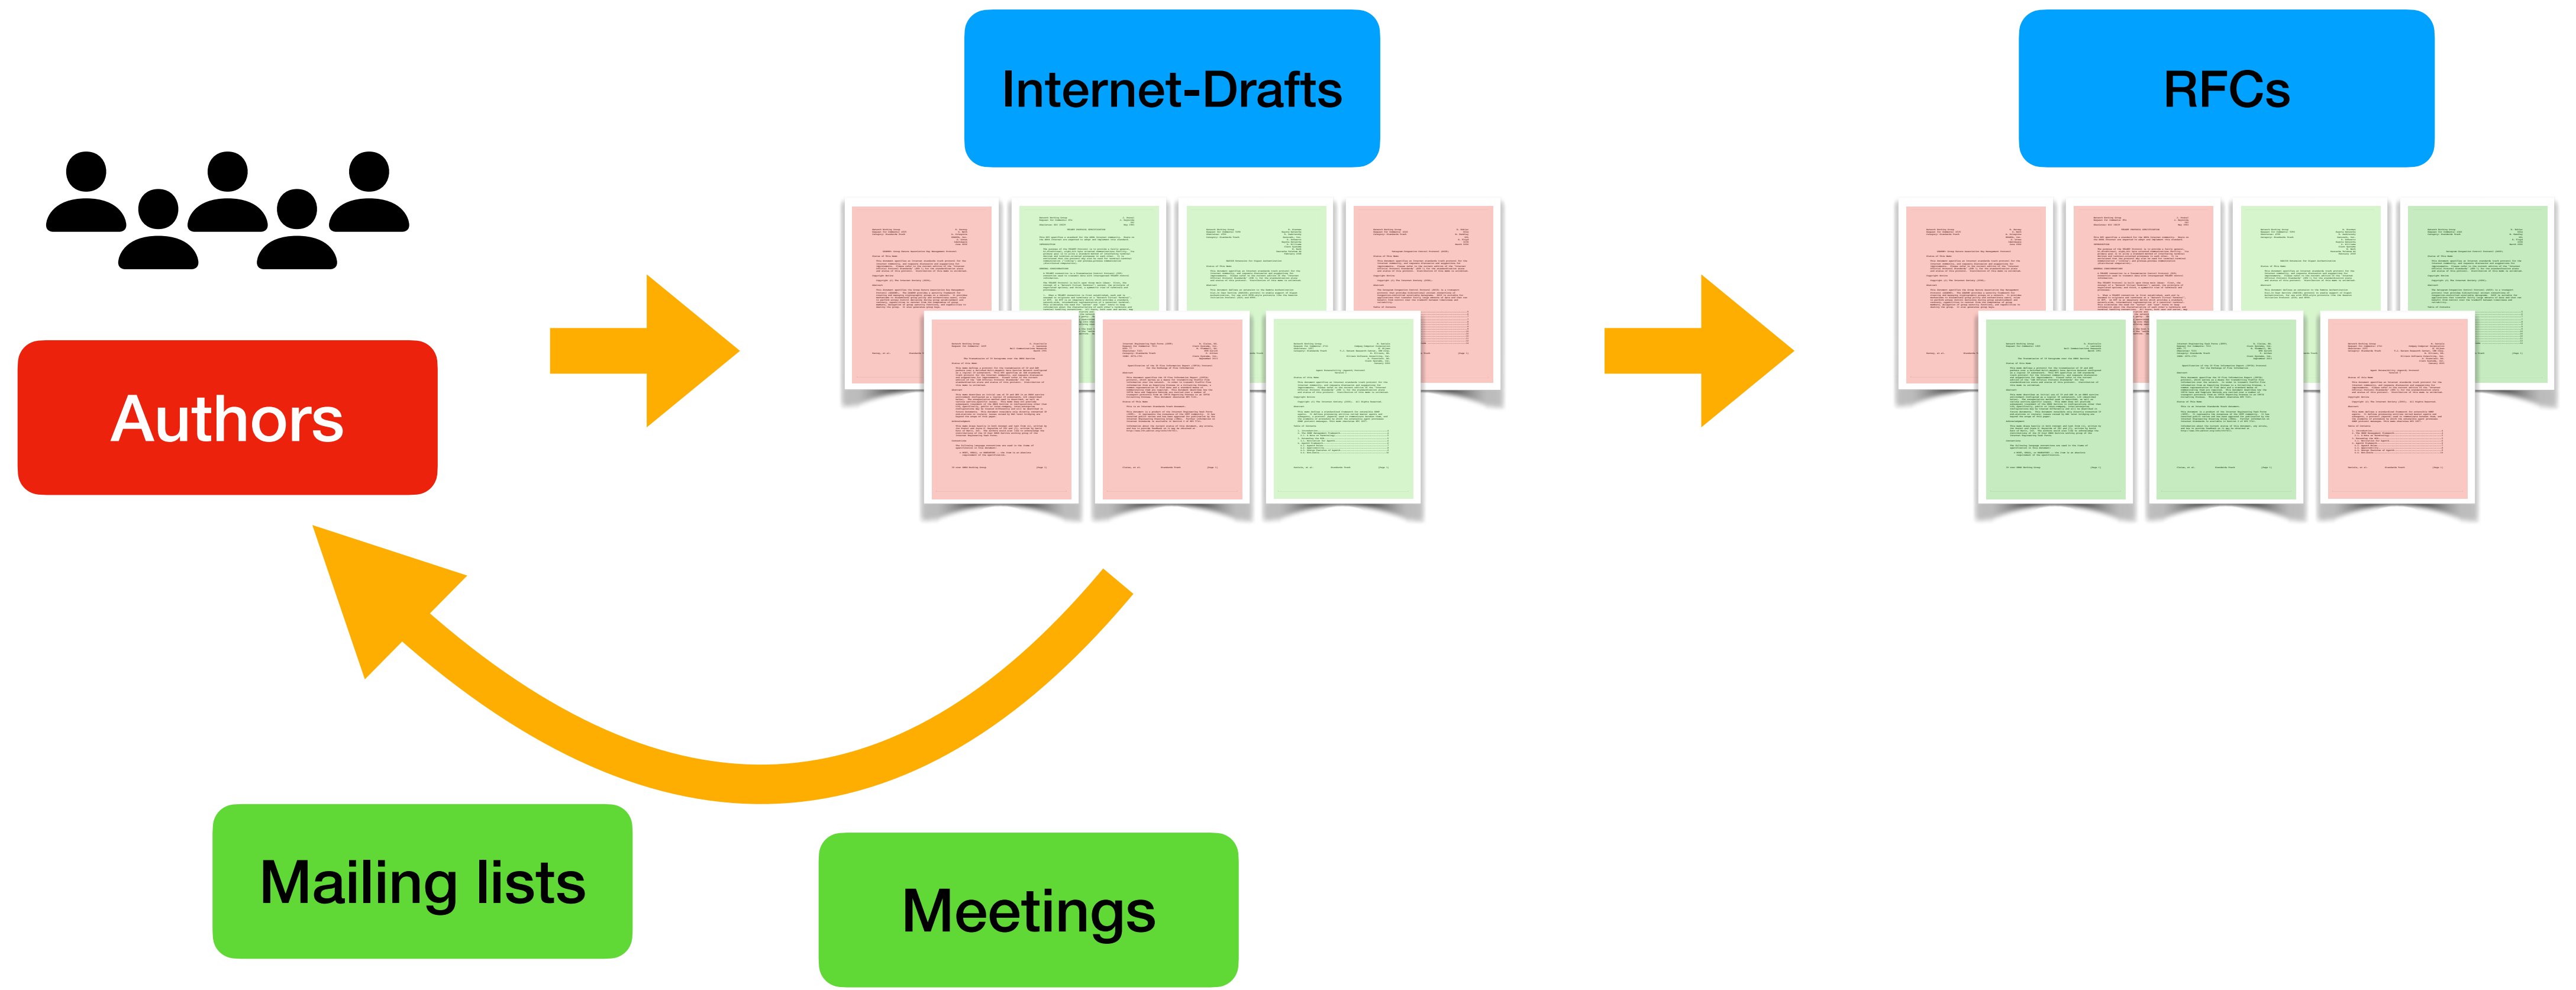




RFCs



























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





Representation from Europe and Asia has grown, with 40% of authors coming from Europe in 2020



Africa and South America are still heavily underrepresented, with only 0.5% of authors coming from either continent in 2020







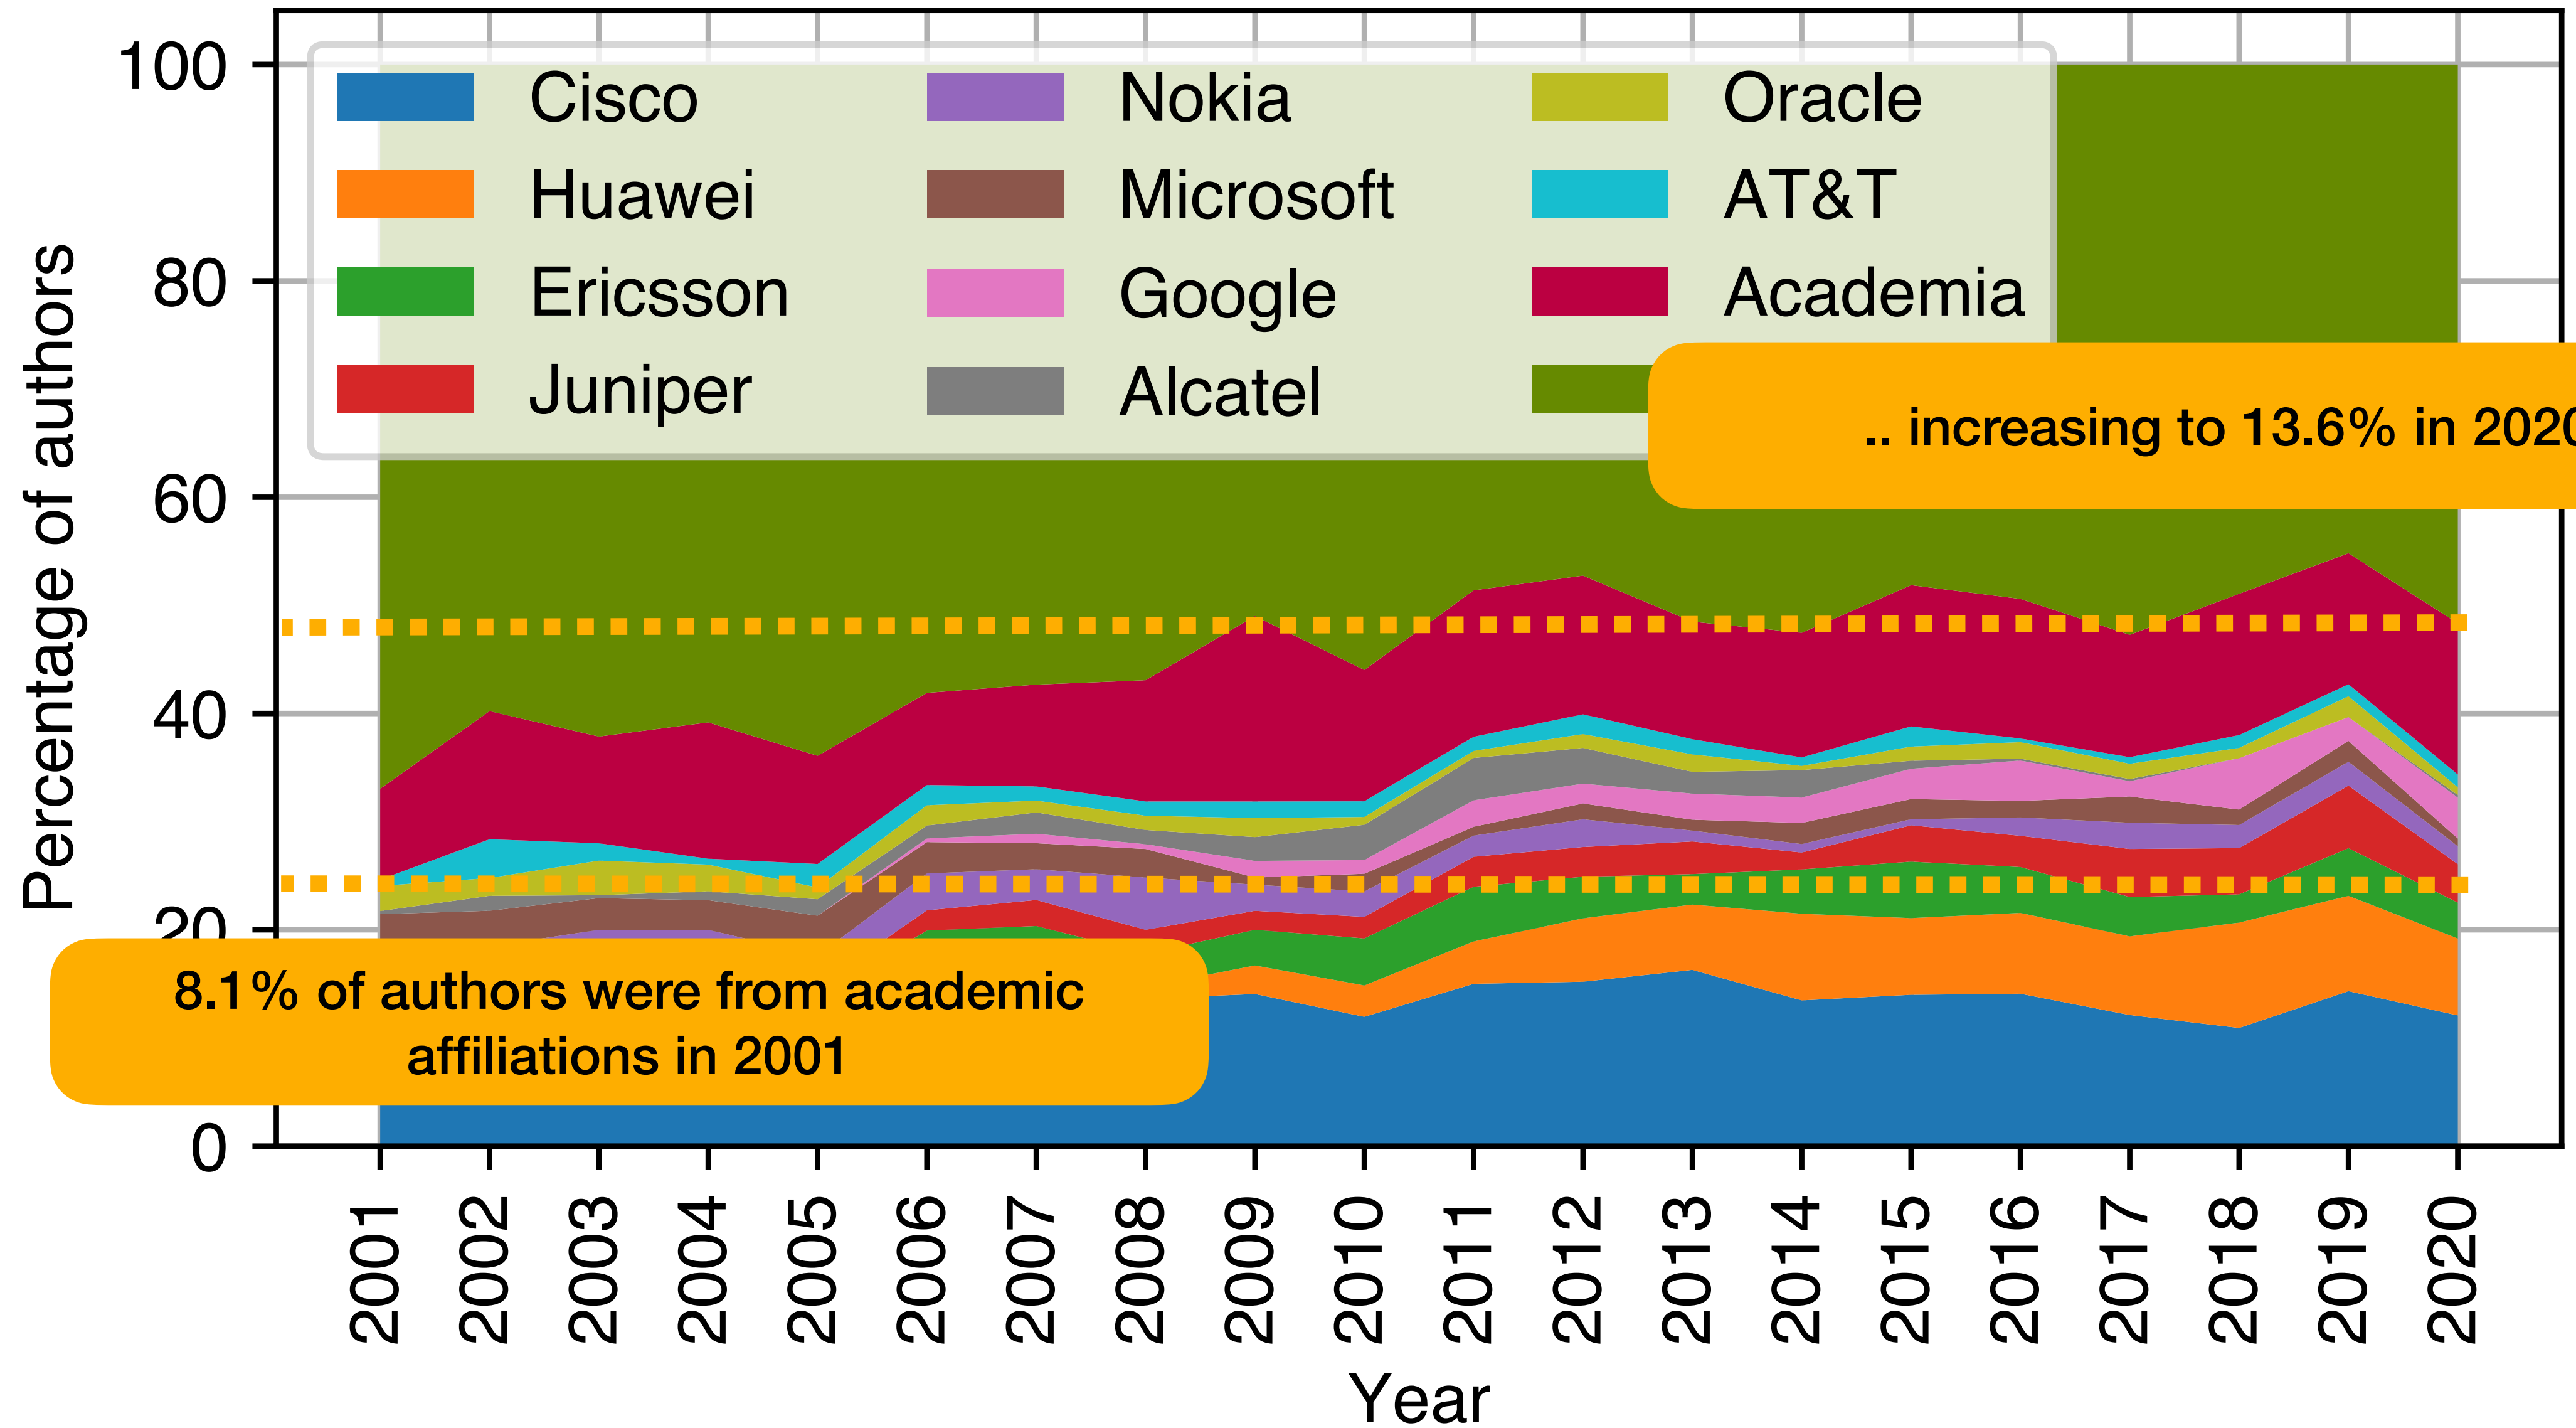


Cisco is the single most frequent affiliation for authors across all years in the dataset













8.1% of authors were from academic affiliations in 2001


.. increasing to 13.6% in 2020






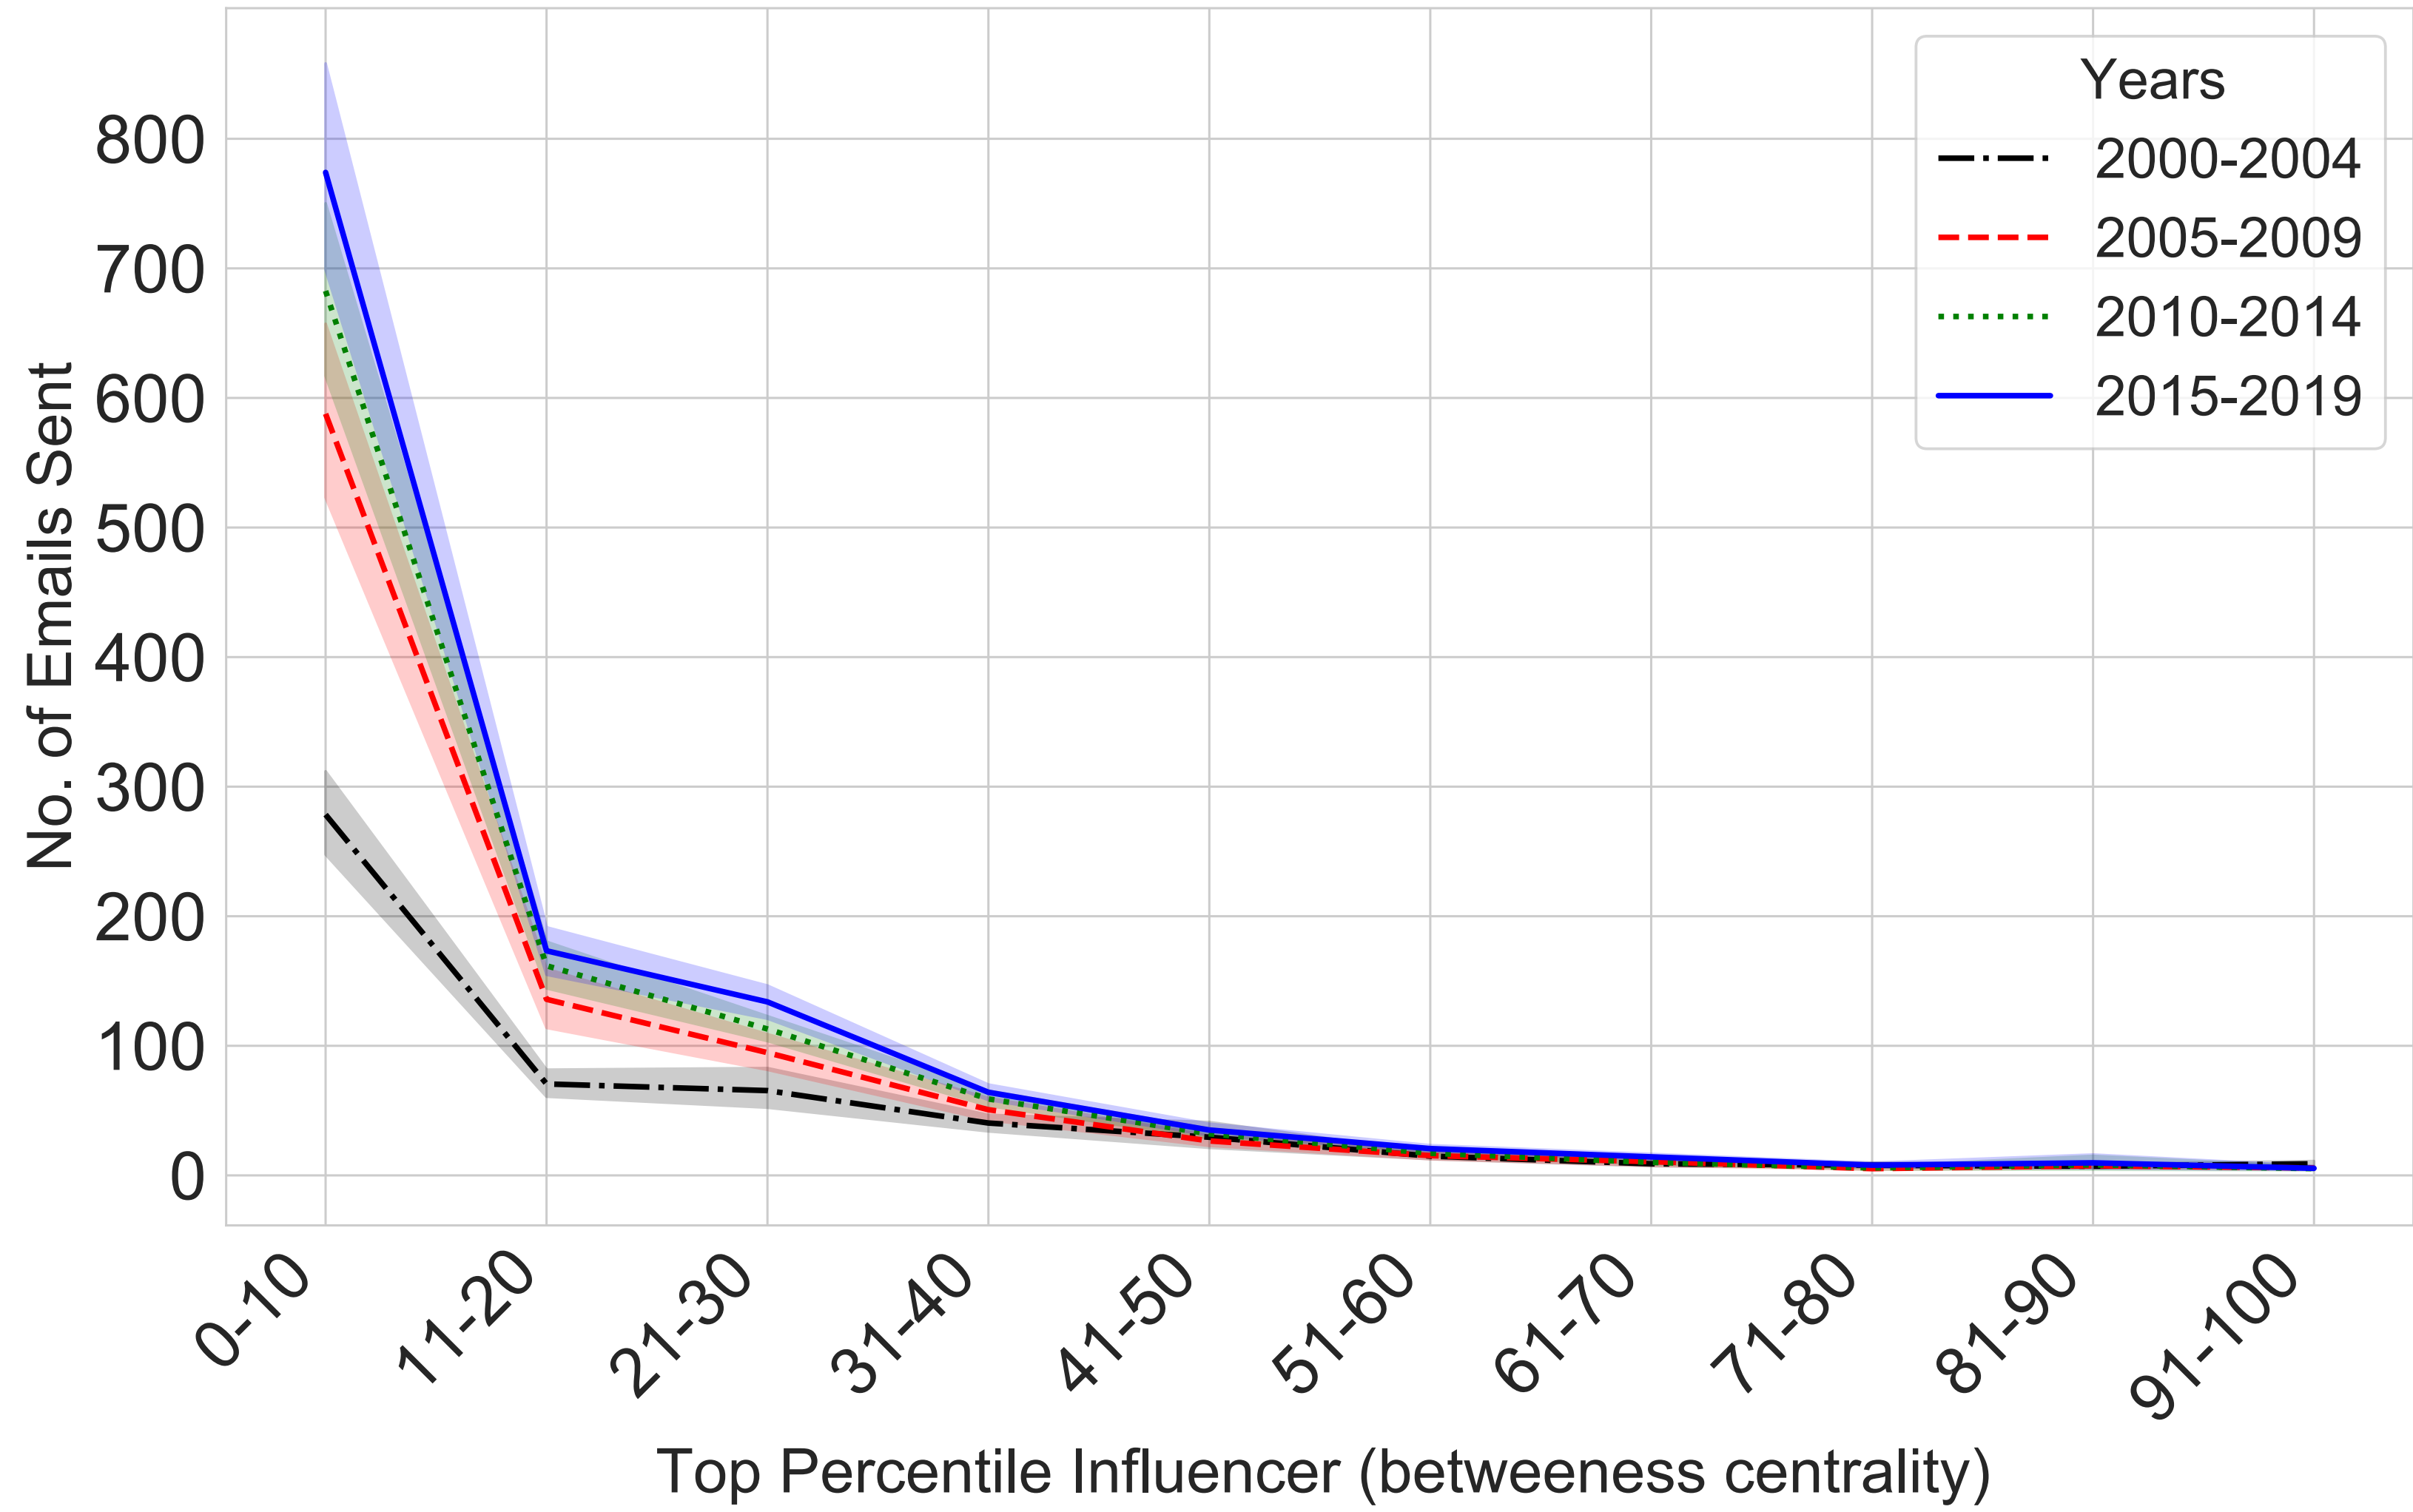


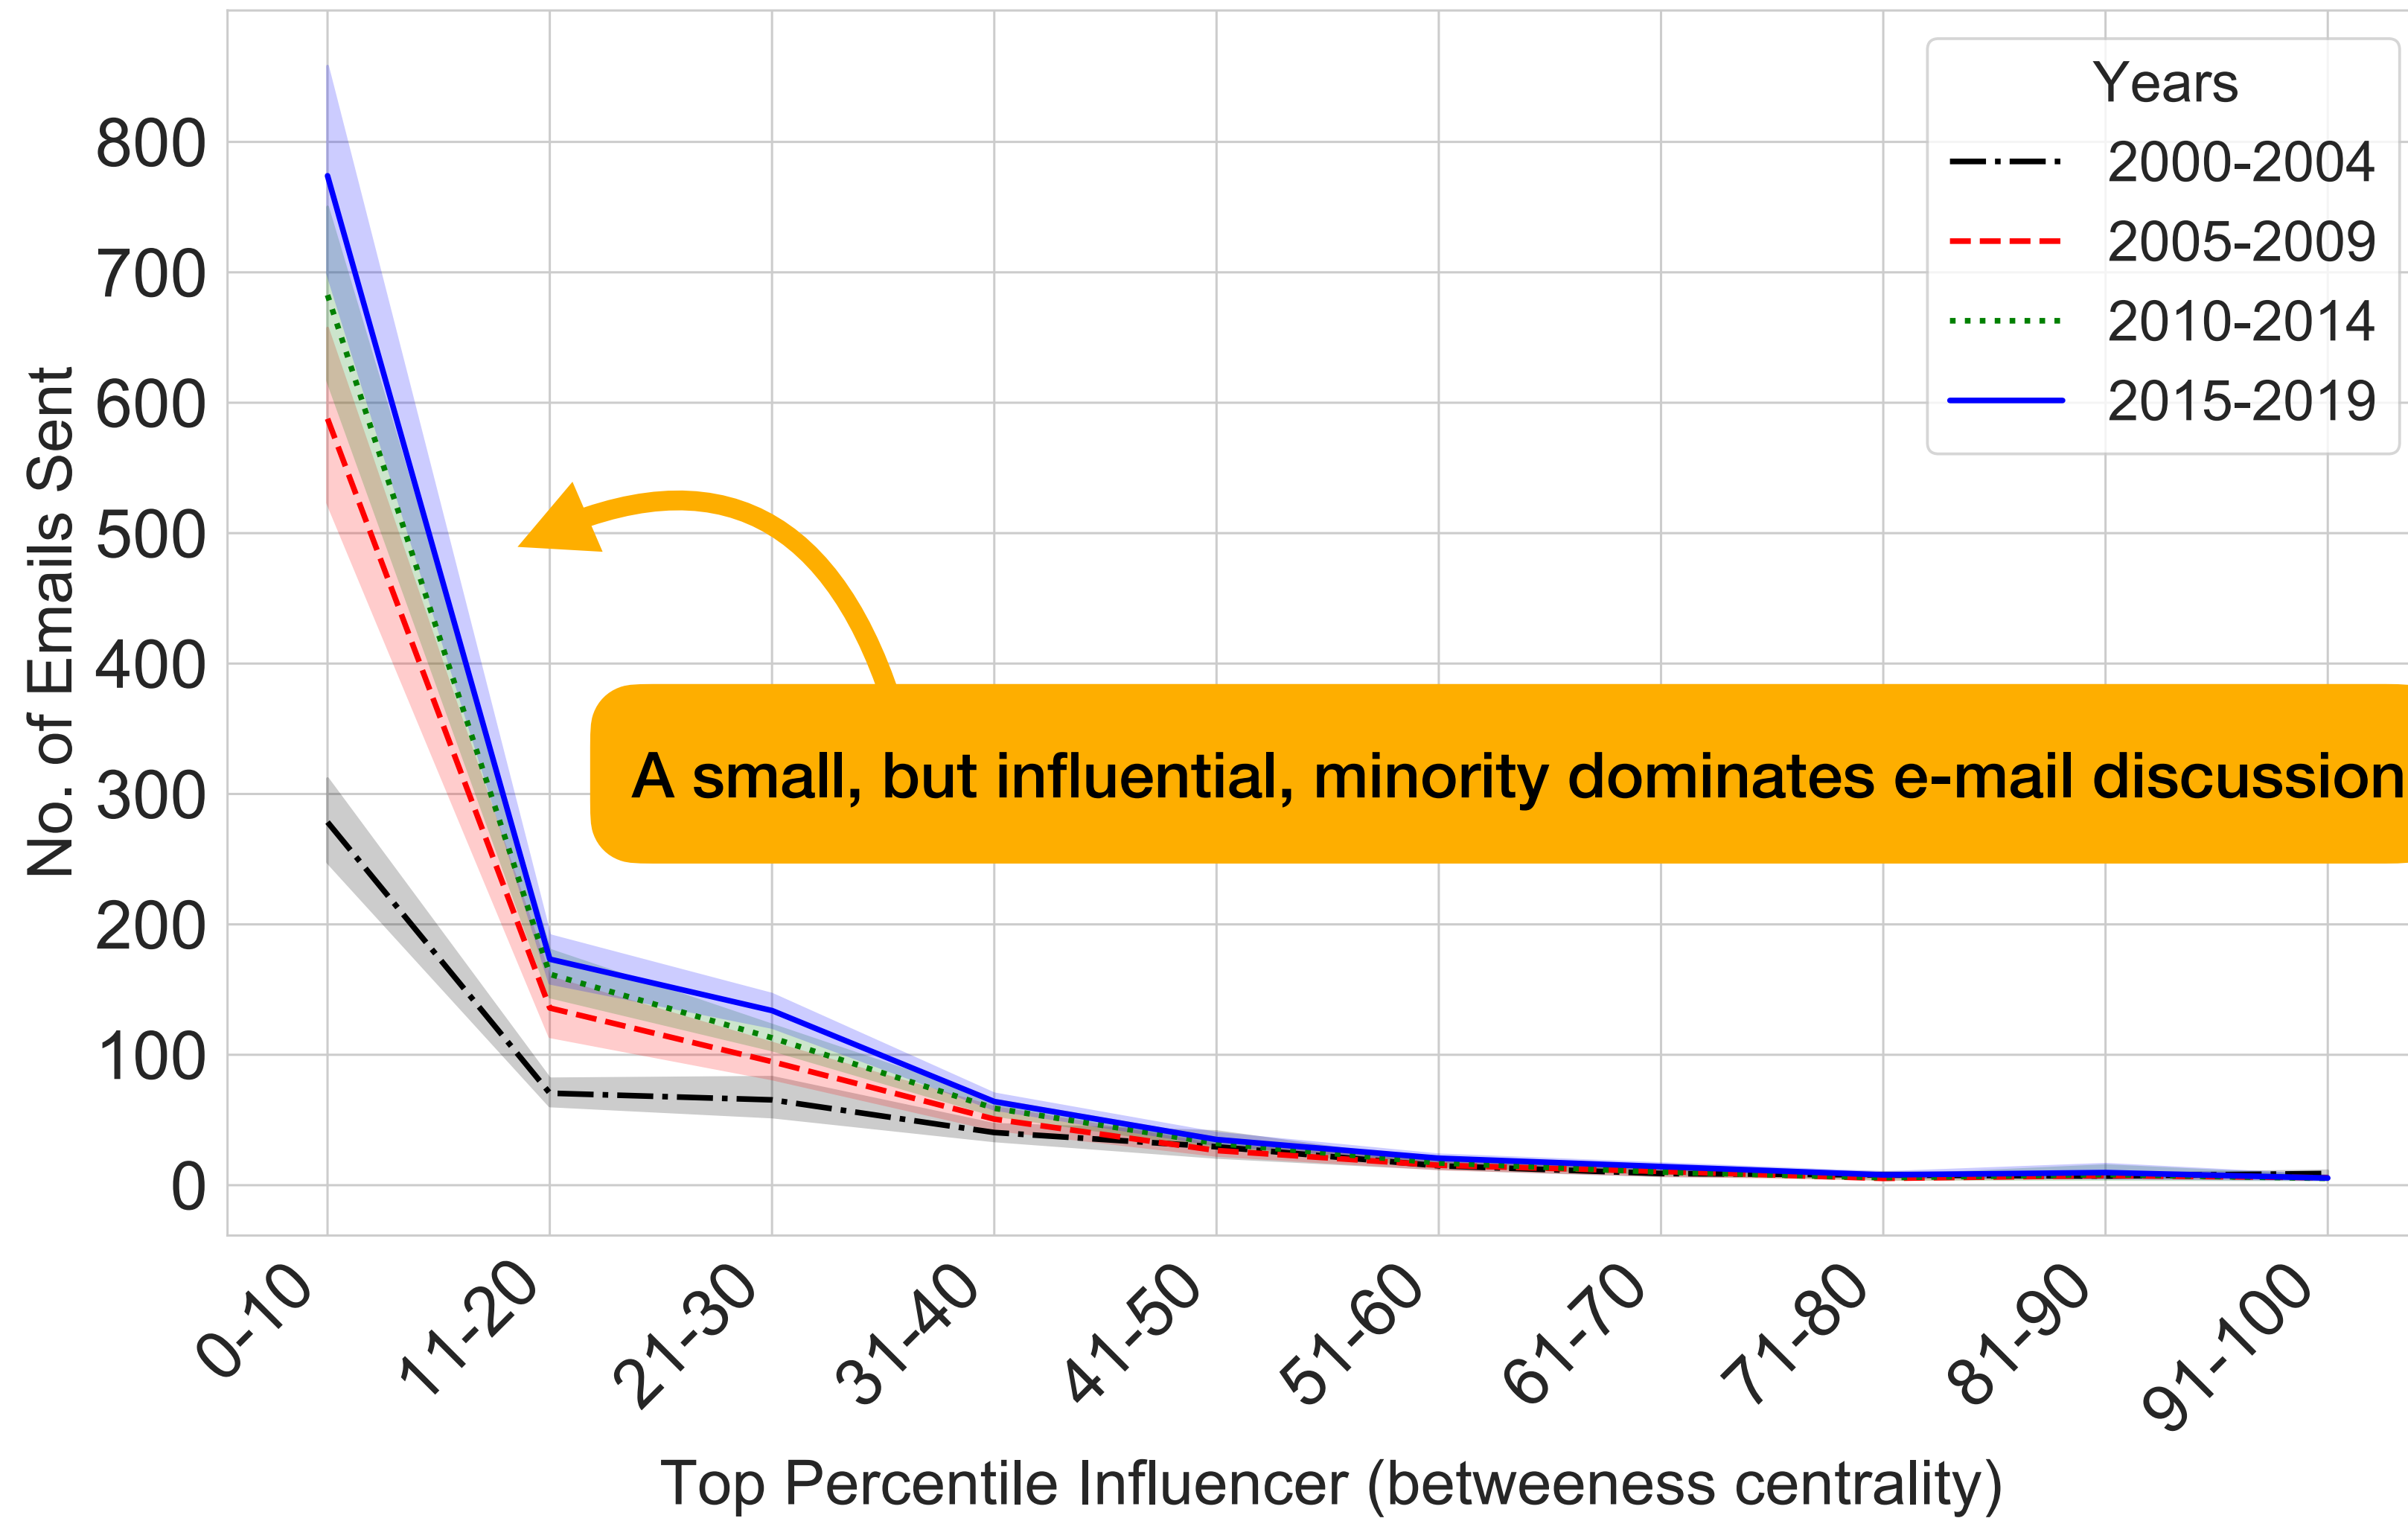





















**RFCs are going through more drafts prior to publication**




Days from first draft to RFC publication

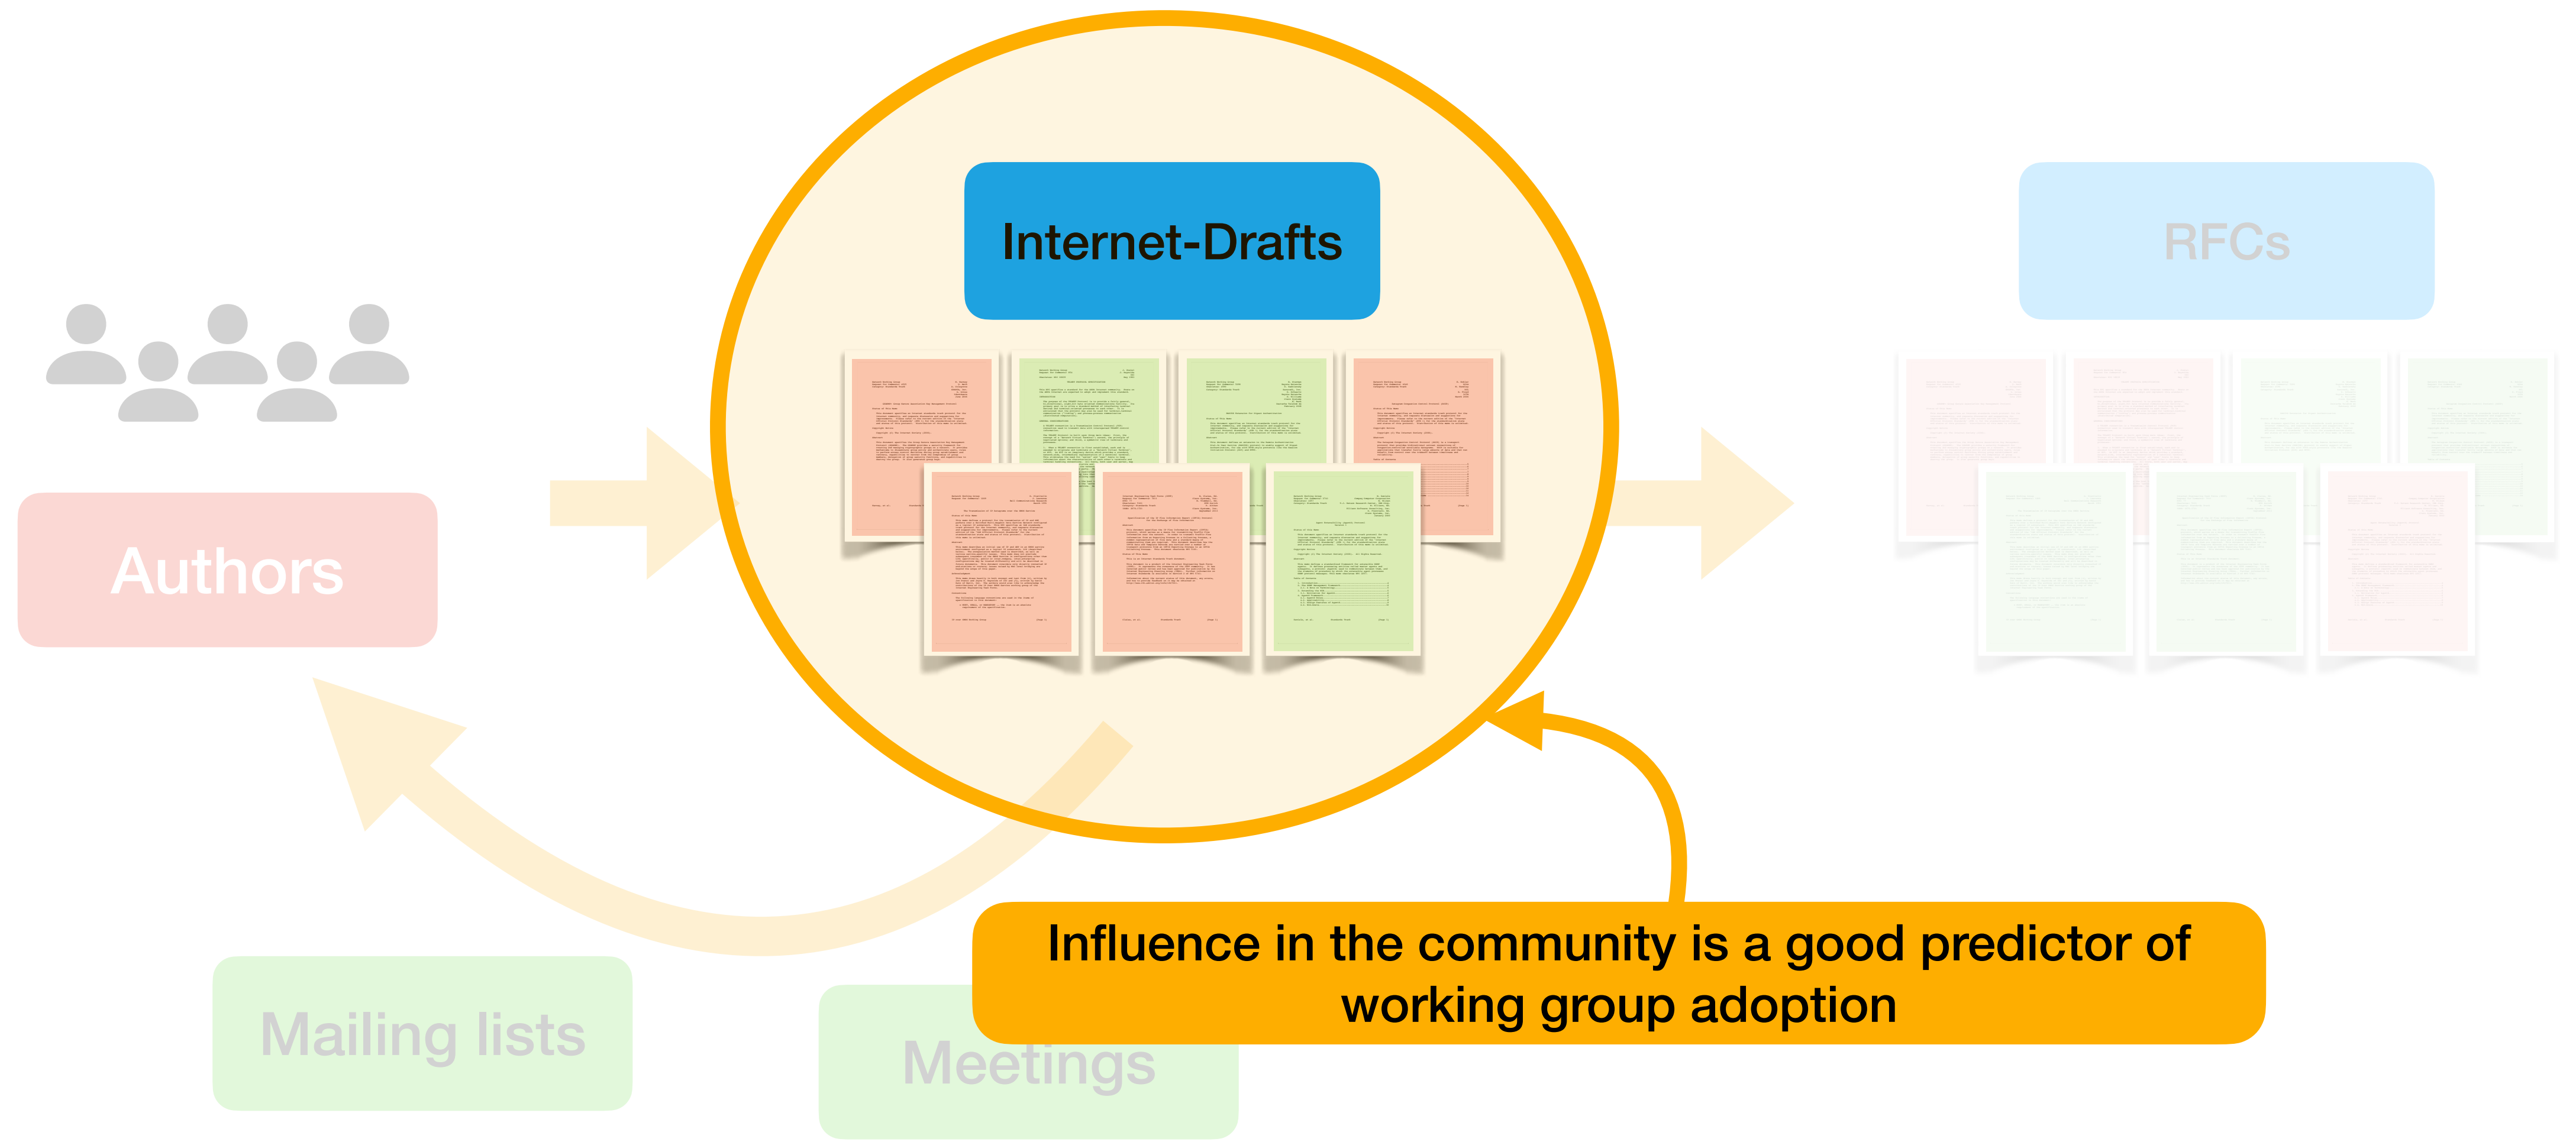


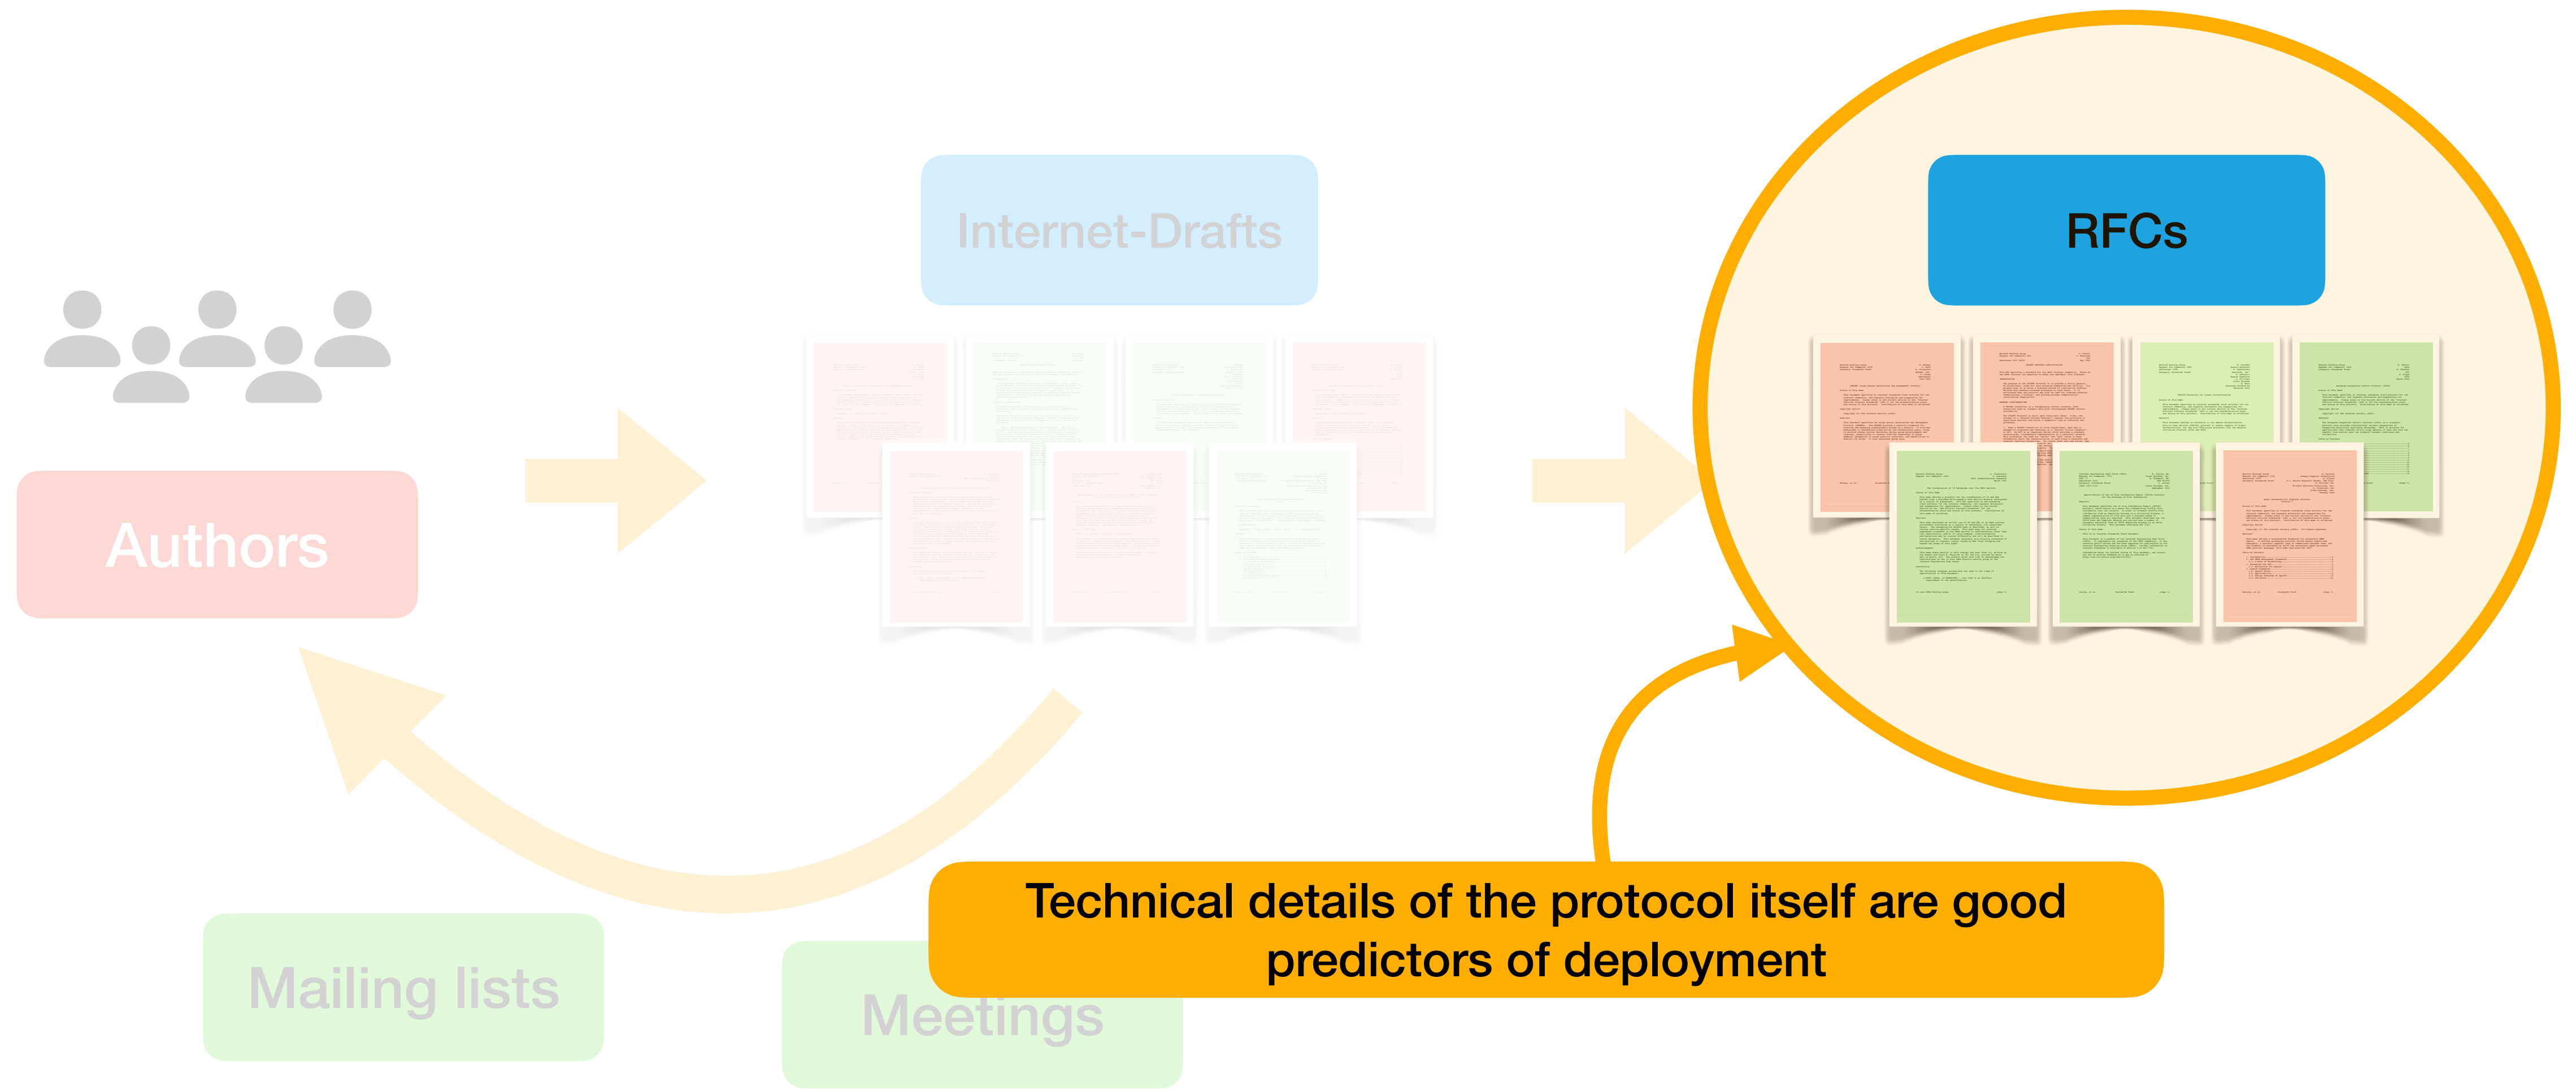
Days from first draft to RFC publication











## Characterising the IETF Through the Lens of RFC Deployment

Stephen McQuistin  
sm@mcquistin.uk  
University of Glasgow

Colin Perkins  
csp@csperkins.org  
University of Glasgow

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

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

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

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

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

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

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

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

### 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>

### 1 INTRODUCTION

Protocol standards are crucial to the successful operation of the Internet. A successful standard provides a basis for interoperability

IMC '21, November 2–4, 2021, Virtual Event, USA  
© 2021 Association for Computing Machinery.  
This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event, USA*, <https://doi.org/10.1145/3487552.3487821>.

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.

With this growing complexity in-mind, we argue that it is vital to gain a coherent understanding of the activities that take place within the IETF, as well as the key factors that may predict the success of a protocol standard.

## The Web We Weave: Untangling the Social Graph of the IETF

Prashant Khare,<sup>1</sup> Mladen Karan,<sup>1</sup> Stephen McQuistin,<sup>2</sup> Colin Perkins,<sup>2</sup> Gareth Tyson,<sup>1,3</sup> Matthew Purver,<sup>1,4</sup> Patrick Healey,<sup>1</sup> Ignacio Castro<sup>1</sup>

<sup>1</sup> Queen Mary University of London

<sup>2</sup> University of Glasgow

<sup>3</sup> Hong Kong University of Science & Technology

<sup>4</sup> Jožef Stefan Institute

p.khare@qmul.ac.uk, m.karan@qmul.ac.uk, sm@mcquistin.uk, csp@csperkins.org, g.tyson@qmul.ac.uk, m.purver@qmul.ac.uk, p.healey@qmul.ac.uk, i.castro@qmul.ac.uk

### Abstract

The Internet Engineering Task Force (IETF) has developed many of the technical standards that underpin the Internet. The standards development process followed by the IETF is open and consensus-driven, but is inherently both a social and political activity, and latent influential structures might exist within the community. Exploring and understanding these structures is essential to ensuring the IETF's resilience and openness. We use network analysis to explore the social graph of IETF participants, based on public email discussions and co-author relationships, and the influence of key contributors. We show that a small core of participants dominates: the top 10% contribute almost half (43.75%) of the emails and come from a relatively small group of organisations. On the other hand, we also find that influence has become relatively *more* decentralised with time. IETF participants also propose and work on drafts that are either adopted by a working group for further refinement or get rejected at an early stage. Using the social graph features combined with email text features, we perform regression analysis to understand the effect of user influence on the success of new work being adopted by the IETF. Our findings shed useful insights into the behavior of participants across time, correlation between influence and success in draft adoption, and the significance of affiliated organisations in the authorship of drafts.

### 1 Introduction

The global success of the Internet owes much to its open development process, a focus on permissionless innovation, and the ready interoperability enabled by its underpinning technical protocol standards. These standards support interworking between a diverse range of systems implemented by different vendors, and encourage the development of a vibrant, open, ecosystem. Given how crucial the Internet has become, it is, however, vital to understand *who* develops and maintains these standards, as they, and the companies they are affiliated with, have the power to fundamentally shape the Internet.

The technical standards that define the Internet are largely developed and maintained by the Internet Engineering Task Force (IETF). The IETF develops and maintains Internet protocols, including those for internetworking and transport

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(TCP/IP and QUIC), routing (BGP, MPLS), security (TLS), and application protocols such as HTTP and WebRTC. The IETF follows an open, consensus-driven process and does not have a formal membership, thereby posing few barriers to entry. The standards it develops are publicly available at no cost and, more importantly for our purposes, the IETF also makes available its email archives, working documents, meeting minutes, etc., providing transparent access to rich datasets that document decades of activities. This allows us to study the process by which the Internet protocols were developed in unprecedented detail. The IETF data provides a representative use case of large-scale, long-lived, distributed online collaboration, and since the dataset long pre-dates the COVID-19 pandemic by several decades, lets us to generate longitudinal insights and patterns pertaining to our research questions.

Protocol standardisation is an inherently social process. Most day-to-day work happens on public mailing lists, aided by meetings, video conference calls, and open document and code repositories. We are specifically interested in better understanding how influence is distributed across stakeholders and how it might affect the standardisation process. This is of critical societal importance: the IETF has a major impact on global Internet technologies, and understanding the social processes involved would give us insight into not only the driving forces behind standardisation, but also its resilience to the loss of major participants. Thus, we ask the following *research questions*:

- (i) How centralised is the active IETF community, and to what extent is it reliant on a small core of participants?
- (ii) How do the most influential participants behave?
- (iii) How does influence (determined by mailing list participation) relate to wider impacts throughout the IETF?
- (iv) Does the organisational affiliation of participants also influence the innovation (adoption of new work) within IETF?

To answer these questions, we collect public mailing list archives (2000–2019) containing more than 2.1M messages from almost 45K senders. We then generate a social interaction graph from these public mailing lists (§2). We find that, akin to many prior social graph studies (Kourtellis et al. 2013; Weitzel, Quaresma, and de Oliveira 2012), influence

## Characterising the IETF Through the Lens of RFC Deployment

McQuistin et al.

ACM Internet Measurement Conference (IMC) 2021

<https://doi.org/10.1145/3487552.3487821>

## The Web We Weave: Untangling the Social Graph of the IETF

Khare et al.

International AAAI Conference on Web and Social Media (ICWSM) 2022

<https://ojs.aaai.org/index.php/ICWSM/article/view/19310>