RFCs have played a pivotal role in helping to formalise ideas and requirements for much of the Internet's design and engineering. They have facilitated peer review amongst engineers, researchers and computer scientists, which in turn has resulted in specification of key Internet protocols and their behaviours so that developers can implement those protocols in products and services, with a degree of certainty around correctness in design and interoperability between different implementations. Security considerations within RFCs were not present from the outset, but rather, evolved over time as the Internet grew in size and complexity, and as our understanding of security concepts and best practices matured. Arguably, security requirements across the corpus of RFCs (over 8,900 at the time of writing) has been inconsistent, and perhaps attests to how and when we often see security vulnerabilities manifest themselves both in protocol design, and subsequent implementation.

In early 2021, Research Director Matt Lewis of NCC Group (global cyber security and risk mitigation specialists) released research exploring properties of RFCs in terms of security [1], which included analyses on how security is (or isn't) prescribed within RFCs. This was done in order to help understand, how and why security vulnerabilities manifest themselves from design to implementation. The research parsed RFCs, extracting RFC data and metadata into graph databases to explore and query relationships between different properties of RFCs. The ultimate aim of the research was to use any key observations and insights to stimulate further thought and discussion on how and where security improvements could be made to the RFC process, allowing for maximised security assurance at protocol specification and design so as to facilitate security and defence-in-depth. The research showed the value of mining large volumes of data for the purpose of gaining useful insights, and the value of techniques such as graph databases to help cut through the complexities involved with processing and interpreting large volumes of data.

Following publication of NCC Group's research, other interested parties read it and identified commonalities with research performed by Mark McFadden (of Internet Policy Advisors LTD), an expert on the development of global internet addressing standards and policies, and an active contributor to work in the IETF and ICANN. Mark had very similar research goals to NCC Group, and in that endeavour he had performed analysis around RFC3552 (Guidelines for Writing RC Text on Security Considerations). RFC3552 provides guidance to authors in crafting RFC text on Security Considerations. Mark noted that the RFC is more than fifteen years old and with the threat landscape and security ecosystem significantly changed since the RFC was published, RFC3552 is a candidate for update. Mark authored an internet draft proposing that, prior to drafting an update to RFC3552, an examination of recent, published Security Considerations sections be carried out as a baseline for how to improve RFC3552. His draft suggested a methodology for examining Security Considerations sections in published RFCs and the extraction of both quantitative and qualitative information that could inform a revision of the older guidance. It also reported on an experiment involving textual analysis of sixteen years of RFC Security Consideration sections.

Matt and Mark are thus very much aligned on this topic, and between their respective approaches, have already gone some way in seeking to baseline how RFC Security Considerations should be expressed and improved. They are therefore seeking to collaborate further on this topic, which will include even further analysis of empirical evidence that exists within the vast bodies of IETF data. Matt and Mark would welcome participation at the forthcoming workshop on analysing IETF Data (AID), 2021. We propose active contribution by way of presentation of our existing research and insights, and would welcome community engagement and discussion on the topic so as to understand how we can utilise the IETF data for the baselining and improvement of security requirement specification within the RFC process.

[1] https://github.com/nccgroup/RFC-Security-Research

[2] https://datatracker.ietf.org/doc/draft-mcfadden-smart-rfc3552-textual-research/00/?include_text=1