Chapter 3
How an IRS Works:
UK's CleanFeed as a Comparative Model

ABSTRACT

In this chapter, the authors describe in detail the UK's CleanFeed design and the blocking mechanisms that it is using. The description presented is based mainly on two papers by Dr. Clayton of Cambridge Computer Laboratory, while many figures are presented in order for CleanFeed's design to be more understandable to a broader public. The only IT expert (with an academic background) who conducted research on CleanFeed software in technical terms is Dr. Clayton.

INTRODUCTION

Having seen the basic blocking mechanisms used today by ISPs (chapter 1), it is time to describe in detail the UK's CleanFeed design and the blocking mechanisms that it is using. BT has never published CleanFeed's design, nor has it stated such intentions for the future. Moreover, the authors' invitation for an interview with IWF's Chief Executive Peter Robbins (back in 2008) was turned down, stating that *"Regrettably, we do not have the resources to respond to interviews and surveys for research purposes"* (see appendix 11).

On the other hand, the only IT expert (with an academic background) who conducted research on CleanFeed software in technical terms is Dr. Clayton of Cambridge University Computer Laboratory. In his paper *Failures in a Hybrid Content Blocking System* (2008) and his technical report *Anonymity*

DOI: 10.4018/978-1-5225-9973-9.ch003

and Traceability to Cyberspace (2005), the CleanFeed's design is thoroughly discussed. According to him, *"This description is based on several separate accounts and, although it is believed to be substantially correct, it may be inaccurate in some minor details"* and *"Up to a point, BT says that my description is mainly accurate, but not entirely so"* (see appendix 8).

The description presented below is based mainly on the two previously mentioned papers, while many figures are presented in order for CleanFeed's design to be more understandable to a broader public. Clayton's research is the only one focusing on the technical aspects of this IRS to such an extent.

Aims

As it is explained in Chapter 1, the most accurate mechanism for blocking content is the content filtering system, which, however, has a crucial disadvantage: high cost of implementation. On the other hand, there are two other basic mechanisms (packet dropping and DNS poisoning) that are very simple and with low cost of implementation, but at the same time very inaccurate (over-blocking and under-blocking issues, and so forth).

Bearing in mind that CleanFeed software was designed to be implemented in BT's customer network (a vast network with a lot of traffic), it is more than obvious that a purely content filtering system would be too expensive. Studying CleanFeed's design (see the next section), it is easy to understand that BT and the UK Home Office tried to develop an accurate system at a low cost, concluding to a 2-stage system using both packet dropping and content filtering mechanisms.

Design

CleanFeed software is a hybrid content blocking scheme, which means that it consists of two separate stages (see figure 1). In brief, the first stage resembles a packet dropping system except that it does not discard requests, but redirects them to the second stage. In the second stage, a web proxy resembles a content filtering system. DNS poisoning was avoided in both stages in order for CleanFeed software not to affect any protocols (such as email) other than web traffic.