

Rearchitecting Security and Troubleshooting Functions in Routers

(Position Statement for PRESTO 2007)

Z. Morley Mao (zmao@umich.edu)

Problem statement: The current Internet lacks strong guarantees for security and reliability. Many functions to ensure security and reliability have been attempted at edge networks or at the application layer, making use of limited support inside the networks. Oftentimes, the lack of sufficient network support severely limits the success of such efforts. For example, troubleshooting tools often rely on traceroute or ping probes and subsequently use inferred, likely inaccurate IP-to-AS mappings to identify responsible networks. Similarly, various routing anomaly detection schemes have been developed relying on routing data from a limited number of vantage points, (*e.g.*, RouteViews and RIPE). From the perspective of edge networks or end users, it is usually impossible to obtain the “ground truths” of detected problems for validation and thus difficult to improve these techniques.

Proposed solution: We examine techniques to enhance security and reliability properties through adding relevant functions inside routers for the core Internet by taking advantage of programmable router platforms. In particular, we try to identify simple router primitives whose compositions enable essential operational tasks for network management. Given the complexity of the existing router software and hardware architecture, we attempt to answer the questions of which functionalities should reside inside routers, and which functionalities should be developed outside the routers in a management layer to improve reliability. This applies to both newly proposed functionalities as well as existing ones. We summarize below the modifications proposed to router functions.

- *Sharing relevant information in response to queries:* As described above, both the static information related to the identity of the router (*e.g.*, the AS to which the router belongs, geographic location) and the dynamic information related to the behavior of the router (*e.g.*, link utilization and packet loss statistics) are critical for network management. Information sharing has direct privacy concerns, as such information exposed may be exploited for malicious purposes. We study how access control can be imposed to eliminate such vulnerabilities. Moreover, to prevent unnecessary overhead imposed by responding to queries for such information, rate limiting is in place, similar to ICMP reply rate limiting. The dynamic information logged include various error messages currently recorded as syslogs, which currently are difficult to interpret due to enormous amount of information. Instead of exposing all state variables, we propose that the routers keep track of the changes and only report unexpected events.
- *Built-in programmable functions:* Related to identifying anomalous events, a desirable function for routers to have is the ability to report deviations from normal behavior, including symptoms of failures and network attacks. The definition of what is considered unusual is not easy to define and needs to be empirically determined and adjusted. In general, we propose to take advantage of programmable router platforms for including functionalities that require dynamic reconfiguration based on external information. For example, traffic filtering or priority-based packet processing capabilities require signature information that is communicated to routers on the fly. Similarly, routers can be conceivably used to make more complex routing decisions based on information within packets in addition to the destination IP address alone, such policies information may need to be dynamically modified using control plane signaling to reprogram the router.
- *Cross layer correlation, layer-independent management:* Routers have access to tremendous amount of useful information useful for network management across many layers including the physical, link, network, and even application layer for unencrypted traffic. We propose that when processing traffic, aside from performing the basic forwarding functions, a router may use potentially idle processors in a multiprocessor architecture for performing other important data correlation, essential for ensuring security and reliability properties. For examples, to ensure forwarding integrity, *i.e.*, the forwarding path conforms to the route advertised, a router needs to correlate between the control plane information with the data plane information across multiple routers. The next hop neighbor router may not follow the advertised path due to routing anomalies or malicious intent. By querying the forwarding data path, a router may verify the forwarding integrity. Although the correlation is performed

on information across layers, the management of individual layers and components is performed independently reduce dependencies.

- *Intelligent management conflict resolution:* One key challenge in network management to carry out various tasks is to identify subtle dependencies and to avoid conflicts across tasks. Conflicts can cause one task to undo the operation carried out by another task. Most operations (*e.g.*, DDoS attack mitigation, traffic engineering) are carried out fairly independently of each other. Currently, conflicts are not systematically resolved. We propose to develop built-in functions inside routers which are modified by different tasks for detecting potential conflicts using internal invariants added. Such invariants can be resource-based, *e.g.*, the link utilization must be below some threshold.