

# Handling Failures and DOS Attacks Using Network Device Groups

Ramkumar Chinchani, Suranjan Pramanik, Ashish Garg  
Dept. of Computer Science and Engineering  
University at Buffalo, SUNY  
Buffalo, NY 14260  
Email: {rc27, pramanik, ashish}@cse.buffalo.edu

*Abstract— With the growing popularity of the Internet and the falling prices of network devices (network interface cards), it is not unusual to find multiple network devices in a computer system. Technologies such as Internet connection sharing and network address translation (NAT) are commonly being used by end users to make network connectivity more viable. In this work, we point out that this implicit redundancy can be used to achieve fault tolerance. It is known that network devices can be grouped to achieve failover support for device failures. In the context of the Internet, security against denial-of-service (DOS) attacks also becomes an important issue. While the use of multiple network devices provides a good solution for device failure, it doesn't guarantee a good defense against DOS attacks. We show that computer systems can become tolerant to DOS attacks if some external factors are considered for network device grouping.*

## I. INTRODUCTION AND MOTIVATION

Network connectivity has become widely available, making it possible to move a large share of communications and transactions to the Internet. Like any software or hardware component, there is always a possibility of failure at any point. The problem of providing reliability and availability (failover support) can be solved at the software or hardware level. Some protocols such as TCP attempt to provide reliable communication over unreliable networks by using packet retransmission. At the hardware level redundancy and replication is used to address fault-tolerance. However, on-board replication is not employed in commonly used network hardware such as PCI or PCMCIA cards because the costs become prohibitive (cost of a chip increases greatly with increase in the surface area). Instead, multiple network interface cards (NIC) are

used to provide failover support. Failures can also happen due to deliberate attacks such as DOS [1] attacks. The idea behind the attack is to overwhelm the target with a large number of service requests with the goal of disrupting availability.

The basic concept of failover switching is not new and it has found applications in the domain of highly available database and computing clusters [2]. In particular, at the network device level, redundancy has been used to rapidly respond to device failures. Intel's Adapter Fault Tolerance (AFT) technology [3] provides continuous network connectivity to high-end servers. SCO MDI driver [4] allows traffic to be shifted to a different configured network adapter card when a hardware failure is detected on the original card. These approaches are very simplistic and are concerned with only host level issues, limiting their effectiveness.

The main focus of this paper is to provide a systematic and comprehensive solution to address the issue of continued network connectivity in the face of both network device failure and DOS attacks. Some outstanding goals and motivation are - (i) *Systematic analysis*: We take into account local factors at the network device level as well as external factors at the TCP/IP protocol and network topology level. (ii) *Rapid and transparent failover switch*: Currently known techniques such as SCTP [5] and "heartbeat" [6] protocols that are used on the top of existing network protocols introduce latency and interfere with rapid failure detection. Instead, we argue that failure detection should be performed at the frontlines of network communication on a host, i.e., at the network device level. (iii) *A cost-effective solution*: Multiple devices are

generally installed on a computer system for the convenience of connectivity, which is utilized to provide cost-effective fault-tolerant network connectivity targeted at end users.

## II. OVERALL TECHNIQUE

A *network device group* is a subset of network interfaces such that each group contains two or more devices. A device failure is detected when errors occur during frame arrival or transmission, and the error code is returned in the status word register of the device. A DOS attack is detected when frame arrival rate is close to the maximum capacity of the network device and the frame arrival rate sustains at this level for at least some preset interval of time that is set by the system administrator. When any of these events occurs, another device in the same group is chosen as a surrogate and all the functions are transparently transferred to this device. The routers on the new path are updated by sending an ARP packet prior to sending the frames. Consequently, upper layers in the network stack are unaware of this switch and doesn't require any state migration.

Not all network devices can be arbitrarily combined into a network device group. If a device in the group should take over as a surrogate for another device, then all the destinations over the network that were reachable from the faulty device should also be reachable from this surrogate device. It is possible to achieve this by finding a common point, typically a router, that is reachable by both the interfaces. For device failures a common point is found very close to both the interfaces so that an ARP update occurs rapidly. In case of DOS attacks, in contrast to device failures, it is desirable to find a common point further away from the network device group since it is necessary that independent paths exist from the two devices to the destination and a DOS attack on one path does not affect the other. These are very important criteria for choosing network devices to form a group.

## III. IMPLEMENTATION

We are implementing the network device grouping inside the Linux kernel (2.4.19). A logical view of the netdevgrp subsystem and its place in the Linux kernel code is shown in Fig. 1. The

entire implementation is divided into two parts: (1) kernel code that responds to the alerts generated by the device driver, and (2) user space code for administration of the network device groupings.

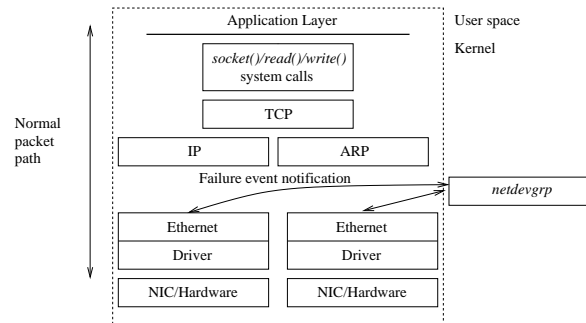


Fig. 1: A logical view of the network

## IV. DISCUSSION

In this paper, we describe a technique that identifies some important network level criteria to perform device grouping to avoid device failures and DOS attacks. Our technique works with device groups that may be heterogeneous such as a combination of wired and wireless network hardware. We have simulated the network device groups in OPNET to verify our hypothesis and study the feasibility. Consequently, we are implementing network device grouping as a separate subsystem in Linux to prove its effectiveness in a lab setting. The prototype can be downloaded from <http://netdevgrp.sf.net>.

## REFERENCES

- [1] L. Garber. Denial-of-Service Attacks Rip the Internet. *Computer*, 33(4):12–17, April 2000.
- [2] Establishing database failover support with HACMP. IBM Websphere Documentation. <http://www-3.ibm.com/software/webervers/appserv/doc/v35/ae/infocenter/was/06061410.html>.
- [3] Intel's Adapter Fault Tolerance Technology. [http://www.intel.com/network/connectivity/resources/technologies/fault\\_tolerance.htm](http://www.intel.com/network/connectivity/resources/technologies/fault_tolerance.htm).
- [4] Failover for MDI devices, 1999. [http://ou800doc.caldera.com/HDK\\_concepts/ddT\\_failover.html](http://ou800doc.caldera.com/HDK_concepts/ddT_failover.html).
- [5] Stream Control Transmission Protocol (SCTP). <http://www.ietf.org/rfc/rfc2960.txt>.
- [6] M. K. Aguilera et al. Using the heartbeat failure detector for quiescent reliable communication and consensus in partitionable networks. *TCS: Theor. Comp. Sc.*, 220(1):3–30, Jun 1999.