Plugging the Leaks Without Unplugging Your Network in the Midst of Disaster

Aaron D Goldman[*], Arif Selcuk Uluagac[†], Raheem Beyah[†], John A Copeland[*]
[*]Communications Systems Center (CSC) [†]Communications Assurance and Performance Group (CAP)
{goldman, selcuk}@gatech.edu {raheem.beyah, john.copeland}@ece.gatech.edu
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332–0250

***Abstract*—Network Disaster Recovery research has examined behavior of networks after disasters with an aim to restoring normal conditions. In addition to probable loss of connectivity, a disaster scenario can also lead to security risks. However, network security has been examined extensively under normal conditions, and not under conditions that ensue after disasters. Therefore, security issues should be addressed during the period of chaos after a disaster, but before operating conditions return to normal. Furthermore, security should be assured, while still allowing access to the network to enable public communication in order to assist in disaster relief efforts. In general, the desire to help with public assistance requires opening up access to the network, while security concerns add pressure to close down or limit access to the network. In this study, we show that the objectives of *availability* and *confidentiality*, two objectives that have not previously been considered together in disaster scenarios, can be simultaneously achieved. For our study, we evaluated six wireless devices with various network configurations, including a laptop, a Kindle Fire e-reader, an Android tablet, a Google Nexus phone, an IP camera, and an Apple TV, to approximate behaviors of a communication network under a disaster scenario. Actual data leakage was tracked and observed for these devices. To the best of our knowledge this has not previously been examined in a systematic manner for post-disaster scenarios. After illustrating the data leakage of various devices, we analyze the risk associated with the various types of leakage. Moving private traffic to a VPN would free the physical network for use as a public resource.**

***Index Terms*—Data Leakage, Network Disaster Recovery, Availability, Confidentiality, Post-Disaster Network Security**

I. INTRODUCTION

The behavior of networks under normal operating conditions has been studied extensively and the security implications of this behavior are relatively well understood. Security practices have concentrated on barring unauthorized parties from access to the network. Unfortunately, security at the network boundary often breaks down under a catastrophic event. Disasters such as fires, floods, hurricanes, tornadoes, explosions, or earthquakes, or terrorist activities, can destroy, not only, the barrier walls that physically enclose networking equipment, but also, networking devices within the facility. This enables unauthorized traffic to enter the network. These realities can cause a network to become vulnerable, from outside or within. Once the intruder is inside, either by physically accessing the equipment or through a new entry point that does not enforce the firewall rules, security protections that guard the exterior of networks are of no benefit.

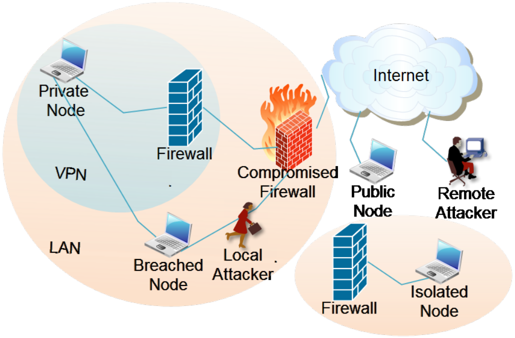

Fig. 1. Network after a Disaster with Remote and Local Attackers

In disaster situations normal operating procedures cease to be in effect. Two examples are the Atlantic hurricane Katrina [1], a natural disaster, which hit the southeastern United States in 2005; and the Russian attack on the nation of Georgia in 2008 [2], a man-made disaster. In both cases, Internet service was interrupted and then restored through extraordinary means, outside the normal security assumptions of communication system designers. In the case of Katrina, connectivity was restored by introducing large numbers of satellite connections. In Georgia, services were restored by moving communication services to new locations outside the disaster area, some of which were hosted by American data centers. This resulted in communication devices connected to unidentified private networks, isolated from the Internet, as well as devices connected to open Internet connections, as illustrated in Figure 1. To study the security implications that occur under disaster scenarios, therefore, we must examine representation of these conditions in the laboratory environment.

The destruction of physical barriers, such as buildings, can subject the network to compromise. There is risk of the network being partitioned (a portion of the network isolated from the rest) or infiltrated (subject to the possible presence of unauthorized equipment), or both. This scenario can easily lead to the intrusion of persons that are not trusted, at a time when authorized persons have evacuated the area. Furthermore, there may be either malicious or opportunistic *cyberlooters* who are trying to steal confidential material. Even more likely, there may be non-malicious actors trying to re-establish communications post-disaster, with or without the permission of the original administrators of the network. In addition to protecting their own information, the administrators may also seek to make the network available for use by others in need of alternative communication channels during a disaster.

The most direct way to address the privacy issue, were that the only goal, would be to ”lockdown” the network so that it refuses to communicate with any device that it can not authenticate. In direct opposition, however, the most useful behavior to aid in disaster recovery effort, were security not an issue, would be to ”open up” use of the network and its Internet connection to the public. These appear to be contradictory goals, however, in our work we address both objectives. As we have stated, under disaster scenarios operations do not proceed as normal. Situations become chaotic. Prior research has concentrated on efforts to re-establish normality. Concerns over network security, however, have not been adequately addressed.

In our study, we examine data leakage that is likely to occur in disaster scenarios. Specifically, we tracked leakage of fourteen configurations of a network with six different connected devices, and examined the leakage from these configurations to determine risk. In Section II, we describe related work that has addressed the behavior of networks in disaster scenarios. The assumptions, scope, and limitations of this work are discussed in Section III. In Section IV, we describe our methodology, consisting of an experimental set up, a traffic analysis, and a risk analysis. Section V presents the results of each of these methodological steps. Conclusions and suggestions for future work are presented in Section VI.

II. RELATED WORK AND MOTIVATION

”Disasters fall into three broad categories: *natural* (floods, earthquakes, and fires), *accidental* (undetected software bugs), and *sabotage* (intentional disaster)” [3] emphasis added. Most of the work related to disaster scenarios has dealt with Network Disaster Recovery (NDR) techniques [4], [5]-[6], [7]-[8]. In general, this area has been explored independently of the research conducted in the area of security. NDR is directed at re-establishing access to the network as efficiently and effectively as possible [9]. Security issues, on the other hand, relate to restricting unauthorized access to the network and, especially in the event of a compromised system because of a disaster, dictate shutting down access to the network. To the best of our knowledge, these two apparently contradictory goals, re-establishing access for availability and restricting access for confidentiality, have not been considered together in disaster scenarios.

The emphasis of the earlier work in NDR was to provide techniques and policies that would provide fast and feasible replacement or repair of communication networks and equipment [4], [10]. They suggest planning for appropriate responses to disasters, but their proposed procedures are activated only after the occurrence of the disaster [5], [6].

In general, networks are subject to two types of failure. First, network connections expected to be functional may have failed. Second, a new connection, which is normally not present, has been added. In the first case, network performance is degraded and applications may not behave as expected. In the second case, unauthorized equipment may gain access to the organization’s network, and network devices may gain unfiltered access to outside networks or the Internet. Moreover, in the midst of, or following a disaster, applications may be running in the background, or may be left running rather than closed on devices that have been abandoned by their normal users. To the best of our knowledge, no previous work has tracked and analyzed the leakage of information that could be accessed by listeners to the network under conditions following a disaster. Hence, in this paper, we systematically analyze the nature of the information that could be disclosed under disaster scenarios, and consider changes to mitigate this security risk.

III. ASSUMPTIONS, SCOPE, AND LIMITATIONS

The data collected in the experiment is observed from a representation that mimics the conditions during a disaster, and not from a real disaster. In the case of an actual disaster, other factors might exist that need to be taken into consideration. One idle device, at a time, is tracked in the study, to provide a reasonable representation of a network system under a disaster scenario. We assume that an attacker is infiltrating the system because of vulnerabilities introduced by the disaster or the NDR techniques. We consider two attack models. The first is of a *local attacker* who gains physical access to the network. The second is of a *remote attacker* who is accessing the network through an Internet connection that is not subject to the firewall rules. These attackers are shown on Figure 1. Furthermore, we assume that the attacker is unable to re-secure the network and therefore, can exploit, but not control the newly created point of infiltration. Attackers with this control are outside the scope of our work.

IV. METHODOLOGY

Our methodology consists of an experimental setup, a traffic analysis, and a risk analysis.

The following experiment was set up to represent configurations found during disasters. The setup was designed to monitor the traffic generated by idle communication devices (e.g., computers, hand-held devices, phones) connected to a network that is open to the public in an attempt to facilitate disaster relief. We connected each device to an unencrypted Linksys wireless access point, recording the traffic between the

TABLE I
CONFIGURATIONS OF NETWORK CONNECTED DEVICES

#	Device	Type	O/S	Configuration
1	Dell Inspiron	Laptop	Linux	Internet (L+I)
2	Dell Inspiron	Laptop	Linux	No Internet (L-I)
3	Dell Inspiron	Laptop	Windows	Internet (W+I)
4	Dell Inspiron	Laptop	Windows	No Internet (W-I)
5	AppleTV	STB	iOS	Internet (TV+I)
6	AppleTV	STB	iOS	No Internet (TV-I)
7	HTC Nexus One	Phone	Andriod	Internet (Ph+I)
8	HTC Nexus One	Phone	Andriod	No Internet (Ph-I)
9	ASUS Transformer	Tablet	Andriod	Internet (T+I)
10	ASUS Transformer	Tablet	Andriod	No Internet (T-I)
11	Amazon Kindle Fire	Tablet	Andriod	Internet (K+I)
12	Amazon Kindle Fire	Tablet	Andriod	No Internet (K-I)
13	D-Link IPcam	Camera	Linux	Internet (C+I)
14	D-Link IPcam	Camera	Linux	No Internet (C-I)

device and the network with WireShark, a network protocol analysis tool. The test was run with six devices in fourteen configurations; each device with (Device+I) and without (Device-I) an Internet connection, as shown in Table I.

The version of Linux used on the laptop was Ubuntu 11.4 with Google's Chrome browser and with the Dropbox client installed. The Windows' setup was Windows 7 with Google Chrome, Firefox, and with the Dropbox client and Microsoft Security Essentials installed. Windows and Ubuntu Linux were chosen due to their popularity in enterprise and consumer usage [11].

During the experiment, the windows laptop was running Internet Explorer, Firefox, and Chrome, each at its new tab page. While running the test on Ubuntu, only Firefox and Chrome were left open to the new tab page. This was done to represent abandoned, but still functioning, connected devices. The applications represent a sample of software likely to be left open during an emergency. Traffic between each device and the access point was recorded for five minutes. We looked at the idle behavior of devices, with no human interference, because users are unlikely to continue operations as usual in a disaster.

Traffic Analysis: We examined the traffic to identify the types of information that are leaking, even when the device appears idle. We inspected the traffic for the presence of known protocols. The packet capture in pcap [12] format was analyzed using WireShark, and protocols of interest were identified for a risk analysis.

Risk Analysis: An analysis was done to identify leakage of confidential or personal information. Each protocol was examined to evaluate the risk associated with it. The identified risks are evidence that additional steps need to be taken to mitigate harm.

V. RESULTS

In this section, we review the results of the traffic analysis and analyze the risks associated with results. These results establish the need for the use of cryptographically secured overlay networks (VPNs) to enable safe communication in the presence of untrusted devices that must be allowed access to the physical network in order to aid in disaster relief. An analysis of the protocols present in the traffic observed under fourteen different configurations are presented in Table II. For more details about these protocols see the specifications at their respective websites [13], [14], [15], [16]. Although the protocols leaked a large quantity of information, much of the information is redundant. Replacing protocols with privacy preserving versions would not be an efficient way to solve the leakage problem. In the long range, better protocols should be written, but in the short range, solutions like the use of VPNs offers a way to protect against improper leakage.

These types of information are highly likely to be disclosed during a period of chaos following a disaster. Some of these protocols, such as ARP, are fundamentally local network protocols and are subject to attack only by local attackers. However, others, such as IE's tracking protection list or Chrome's safe browsing, are Internet protocols. For example the IP Camera sends pictures back to the manufacturer's website unless blocked by a firewall. In the absence of a reliable protection mechanism, (e.g., firewall), these protocols can leak information that will be subject to attack by both local and remote attackers. Information can be used by cyberlooters to build a dossier of the organization. Information about equipment and software can be used to design a targeted technical attack, that could be performed long after conditions have returned to normal. The information about individuals and organizational structure can be used to design socially engineered components of an attack days or even weeks after the disaster. Names given to computers often reflect the names of employees, divisions, or departments. Similarly, names applied to files and shared folders reveal relationships between these entities. Web histories can be used to learn about specific roles or interests of individuals. In combination, and when added to external public data, all of this information can be used or sold by those with malicious intent.

VI. CONCLUSIONS AND FUTURE WORK

One objective of the behavior of communication networks, under disaster scenarios, is to assure the safekeeping of confidential information by closing down or restricting use of the network. This is the confidentiality objective. The other is to allow public access to the network by opening up or expanding access, thereby providing alternatives to normal communication channels which may have been disabled or destroyed [22]. This is the availablity objective. These two objectives had not been considered simultaneously for post-disaster senarios.

In this study, actual data leakage was tracked from fourteen configurations of a network with six different connected devices. We examined the resulting leakage to determine the risk presented. The evidence of leakage from apparently idle devices connected to the network illustrated the need for a security solution that would meet both of the apparently contradictory objectives of availability and confidentiality.

A number of related areas remain for further consideration. We will complete a cost-benefit analysis of possible security mechanisms suitable for post-disaster conditions.

TABLE II
ANALYSIS OF PROTOCOLS LEAKING INFORMATION

PROTOCOL NAME Configuration Number - Table I	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
ARP : Address Resolution Protocol [14]	X	X	X	X		X	X	X	X	X	X		X	X
Translates addresses - MAC to IP Address reveals equipment vendors and locations														
DHCP : Dynamic Host Configuration [14]	X	X	X	X						X		X	X	
Network configuration protocol Reveal IP addresses, ISPs, router information														
DHCPv6 : Dynamic Host Configuration V6 [14]	X	X	X	X										
Network configuration protocol V6 Same as above, IPv6 is enabled, displays neighbor advertisements														
DNS : Domain Name System [14]	X	X	X	X		X	X				X		X	X
Converts domain names to IP addresses Reveals website access attempts, security software version, virus definitions, payload URLs, browsers in use, Chrome Safe Search, IE Tracking Protection Lists, expected computer names														
DROPBOX : Dropbox LAN Sync [17]	X	X	X	X										
Advertises Dropbox instances Reveals presence of dropbox, user ID#, shared folder ID#, organizational structure Pseudonyms can be de-anonymized														
EAPOL : Extensible Authentication [14]												X		
Authentication framework Reveals network key infrastructure														
HTTP : Hypertext Transfer Protocol [16]	X	X	X	X		X	X	X	X		X		X	X
Transfers files and applications across the web Reveals Applications														
ICS lap : Internet Connection Sharing [15]			X											
Legacy proxy configuration protocol Can be used to redirect traffic, reveals Windows and ICSlap														
IGMP : Internet Group Management [14]	X	X	X	X	X	X							X	X
Connects multi-cast subscriptions Reveals subscriptions, organizational structure														
JABBER protocol [18]	X	X	X	X										
Instant messaging protocol Reveals chat networks in use														
LLMNR : Link-Local Multi-cast Name Resolution Protocol [14]			X											
Decentralized computer name lookups Reveals expected computer names, LAN, Name information about users, can be used to de-anonymize Dropbox														
MDNS : Multicast DNS [14]			X	X		X							X	X
Decentralized DNS lookups for LAN Reveals devices and services on LAN														
NBNS : NetBios Name Service [14]	X	X	X	X										
Broadcasts mapping of equipment on network Reveals names, makes, and models of computers and routers														
OCSP : Online Certificate Status [19]	X													
Looks up revocation status of certificates Reveals attempts to access applications														
SANavigator : SAN Navigator [20]			X	X										
Broadcasts network storage equipment information Reveals SAN, folder names														
SOCKS [14]			X											
Internet web proxies looks ups Reveals the Internet usage policy														
SSDP : Simple Service Discovery [14]			X	X									X	X
Advertises services available on network Reveals network services, printers, firewall bypass, media servers, active equipment, and content directories, schema URLs for unrecognized services														
XMPP : Extensible Messaging Presence [21]	X	X	X	X										
Communicate with numerous Google services Reveals use of Google services														

ACKNOWLEDGMENTS

We thank Sakthi Vignesh Radhakrishnan of the Communications Assurance and Performance Group at Georgia Institute of Technology for set up and operation of the test bench.

REFERENCES

[1] Alin Popescu James Cowie and Todd Underwood. Impact of Hurricane Katrina on Internet Infrastructure. September 2005.

[2] Stefanie Hoffman. Russian Cyber Attacks Shut Down Georgian Websites. *CRN.com*, August 2008.

[3] G.P. O'Reilly, D.J. Houck, E. Kim, T.B. Morawski, D.D. Picklesimer, and H. Uzunalioglu. Infrastructure simulations of disaster scenarios. In *Telecommunications Network Strategy and Planning Symposium. NETWORKS, 11th International*, pages 205–210, June 2004.

[4] Chi-Ming Chen, A. Macwan, and J. Rupe. Network disaster recovery [Guest Editorial]. *IEEE Communications Magazine*, 49(1):26–27, January 2011.

[5] R.E. Krock. Lack of emergency recovery planning is a disaster waiting to happen. *IEEE Communications Magazine*, 49(1):48–51, January 2011.

[6] J.C. Oberg, A.G. Whitt, and R.M. Mills. Disasters will happen - are you ready? *IEEE Communications Magazine*, 49(1):36–42, January 2011.

[7] M. A. Pimentel-Nino and M. A. Vazquez-Castro. Cross layer content delivery optimization for holistic network design in disaster preparedness and recovery scenarios. In *IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6, July 2011.

[8] K. Takahata, S. Takada, and Y. Shibata. Disaster communication network by combination of different wireless lans. In *AINAW 22nd International Conference on Advanced Information Networking and Applications - Workshops*, pages 1129–1133, March 2008.

[9] K.T. Morrison. Rapidly recovering from the catastrophic loss of a major telecommunications office. *IEEE Communications Magazine*, 49(1):28–35, January 2011.

[10] Yang Ran. Considerations and suggestions on improvement of communication network disaster countermeasures after the wenchuan earthquake. *IEEE Communications Magazine*, 49(1):44–47, January 2011.

[11] Distro Watch. *http://distrowatch.com/*.

[12] libpcap. *http://www.tcpdump.org/pcap.html*.

[13] IEEE standards. *standards.ieee.org*.

[14] The Internet Engineering Task Force. *ietf.org*.

[15] Microsoft open specifications. *www.microsoft.com/openspecifications*.

[16] World wide web consortium (w3c) standards. *www.W3.org*.

[17] Dropbox. *www.dropbox.com*.

[18] Jabber.org. *www.jabber.org*.

[19] International Telecommunication Union. *www.itu.int*.

[20] sanavigator. *www.sanavigator.com*.

[21] XMPP Standards Foundation. *xmpp.org*.

[22] K. Mase. How to deliver your message from/to a disaster area. *IEEE Communications Magazine*, 49(1):52–57, January 2011.