Penetration Analysis of a XEROX Docucenter DC 230ST: Assessing the Security of a Multi-purpose Office Machine *

Thomas E. Daniels, Benjamin A. Kuperman, Eugene H. Spafford

CERIAS T.R. No. 99-09

Center for Education and Research in
Information Assurance and Security (CERIAS)

1315 Recitation Building
Purdue University
West Lafayette, IN 47907–1315
{daniels,kuperman,spaf}@cerias.purdue.edu

1 Executive Summary

Recent advances in manufacturing technology have made possible multi-purpose office equipment that handle a large variety of tasks previously relegated to multiple individual machines. In this paper, we examined the Xerox Docucenter 230ST (DC 230ST) which supports copying, faxing (both sending and receiving), network printing, and web-based printing, among other features.

Because of the shared comnunication medium, machines that physically host multiple services may pose a greater security risk than individual devices. In the case of the DC 230ST, a CPU and hard drive control the functions of the of the above features.

We found that if an attacker can gain physical access to the machine, the programming of the machine can be compromised. A sophisticated attacker can subvert the machine without resorting to opening the physical casing. We were unable to compromise the machine remotely, although we did not exhaust the possibility of such a compromise. The results of our attempts along with recommendations regarding the possible deployment of a DC 230ST are contained within this document.

*Portions of this work were supported by the sponsors of CERIAS.

2 Introduction

2.1 Overview of the Problem Space

Over the past few years, many office equipment manufacturers have been marketing multipurpose office machines that combine the functionalities of many single purpose machines. Frequently, these machines combine photo-static reproduction ("copying") with faxing capabilities. Recent changes have led to network printing features as well. This approach eliminates the redundancy of having multiple devices, each of which has a similar printing mechanism. With the addition of the image capture of a photocopier and the network connection from the printer, adding scanning capability only requires the addition of memory and some temporary storage for the image. The result is a specialized computer that has a built in scanner, fax modem, network card, disk drive, RAM, and laser printer all within a single casing. For an office environment, this allows a single piece of hardware to handle the functions of four common pieces of office equipment at a fraction of the cost[1].

However, this combination of services introduces new potential security threats by means of **cross-channel communications**, or communications between multiple services [TDS99] that previously were not possible. One example would be an attacker configuring a machine to store everything submitted to the printer via the network card so that the stored information could later be faxed to the attacker. In this paper we discuss our examination of one such device for potential threats including cross-channel communications.

2.2 General approach

To examine the potential vulnerabilities we:

1. Obtained a multi-purpose office center device.
2. Examined the vulnerabilities of each of the services that it offered.
3. Examined the new vulnerabilities introduced by this combination of services.

3 About the Xerox Document Center 230ST

For this project, we were supplied a Xerox Document Center 230ST (DC 230ST). The DC 230ST is a multi-function copier, with the following features:

- Laser printer print technology
  - Postscript
  - PCL5
- Copier

[1] This also means that there is a single point of failure for many frequently used services

    - Multiple paper/transparencies trays
    - Collator
    - Stapler
    - etc.
- Network Interfaces
    - Ethernet
    - Token Ring (not included in our Docucenter)
    - Parallel Port
- Network Protocols
    - Novell Netware
    - TCP/IP
    - AppleTalk
    - Banyan Vines
    - NetBIOS/NetBEUI
    - NetBIOS/IP
- Network Services
    - Line Printer Service (LPR)
    - HTTP
- Fax
- Network Fax
- Print from floppy disk
- Scan to File (not installed)
- Fax Mailboxes (not installed)
- Remote administration (not installed)

In addition to those previously mentioned, the Docucenter also has a few undocumented physical interfaces of interest. There are two 9 pin serial interfaces on the Docucenter. One is directly under the user control panel. The second is on the back of the unit. There are also some connectors of unknown type on the rear of the copying unit. They are likely used to connect additional components to the Docucenter.

The DC 230ST appears to be composed of two semi-independent control systems. The copier control unit seems to control the faxing and copying functionality. The internal computer handles network services and print rendering. This was observed when the internal computer crashed, disabled network services, but still allowed the user to copy documents. Exactly how the two different control systems interface with each other is unknown, but the answer may have strong implications for cross-channel communication.

4 Our Approach

In this section, we present an overview of our testing methodology and the subsystems that we tested.

4.1 Testing Methodology

Our approach to testing the DC 230ST for security problems was derived from Richard Linde's seminal paper "Operating System Penetration" [Lin75]. Based on Linde's work, we break our methodology into the following four stages.

- Knowledge of the system control structure
- Generation of flaw hypotheses
- Testing the flaw hypotheses
- Generalization of discovered flaws into general system weaknesses

Knowledge of the system control structure is important because one must understand the test system in order to find its flaws. This knowledge is especially important for testing the DC 230ST because it is not a conventional computing system; it may be prudent to consider flaws or policies that do not apply to a more traditional computing system.

When the system is understood to some extent, we generate flaw hypotheses. A flaw hypothesis is an "educated guess" that a flaw exists in a system. It may be based on flaws documented in other systems or because the test system acts in unexpected ways under some conditions. An example of a flaw hypothesis that we generated after some experience with the DC 230ST is "No IP-based access control for the HTTP service exists on the DC 230ST." The hypothesis should be concise, clear, and testable.

Testing of flaw hypotheses involves designing an experiment to test each hypothesis. The experiment should include the hypothesis, a detailed procedure to be followed, and criteria to evaluate whether the flaw exists based upon the outcome of the experiment. For the example hypothesis given above, the experiment would be a thorough examination of the user interfaces and documentation of the DC 230ST for reference to any access control mechanism for the HTTP interface. In this case, interpretation of the results is simple in that if an access control mechanism is found, the flaw is not there. If an undocumented feature cannot be found after a thorough search for it, we argue that the actual presence of the feature, however unaccessible, should be treated as if it did not exist.

We generalize flaws into system weaknesses by explaining or showing how the flaw can be used in a certain environment to compromise the security of the device. In our example, the flaw generalizes in several ways. First, none of the IP-based network services have an access control mechanism. Secondly, if the system is installed on a widely accessible network without a firewall or other packet filter, and the HTTP service is enabled, then anyone on that network can monitor the jobs on the device and submit a print job which may include postscript viruses, etc.

It is important to note that these different stages are not serial in nature. In some ways, each may occur simultaneously during active testing. For instance, it is common during testing to discover unusual system behavior unrelated to the testing in progress. The discovery adds to your system knowledge and may prompt the generation of additional flaw hypotheses, thereby leading to new experiments.

4.2 Areas of Testing

We divided our testing into the following three areas.

- Telephone/Facsimile
- Network
- Physical

A test is categorized into one of the three categories if the test data sent to the DC 230ST is a logical component of that category.

There are some notable areas that we left untested. We did not test the DC 230ST for electromagnetic emanation of information because we do not have the required facilities for this type of testing. Similarly, we did not consider possible information leakage via power line surges during usage.

5 Summary of Results

5.1 Telephone/Facsimile

We attempted to subvert the Docucenter via its telephone interface by dialing into it using a conventional data modem. Despite attempts at a wide variety of baud rates and data settings, we were unable to form a connection with the DC 230ST. The documentation refers to a remote management option that is available for the Docucenter, but this option was not included on our unit.

By setting the same data modem to FAX mode, we were able to connect and send faxes using a freely available FAX program called sendfax [Doe]. One high level feature of the fax protocol allows passing a `station id` from the sender to the receiver. It is the `station id` that is often printed along the top of received faxes. We attempted to modify sendfax to send an abnormally long `station id` to the DC 230ST during fax transmittal. Unfortunately, our modem would not accept the long `station id` therefore the results are inconclusive.

We also tried sending pseudo-random strings of data instead of valid FAX messages, but the Docucenter caught these as illegal facsimiles and disregarded them cleanly.

A fairly comprehensive overview of the threats affecting Fax machines is presented in [OL95]. Many of the attacks/threats described required the services of an "*electronics hobbyist*" to which we did not have access.

5.2 Network

There are two basic levels of the DC 230ST's network implementation into which we break our testing: low level tests and high level tests. Low level tests focus on vulnerabilities in the network and transport layers of the network protocol implementation such as

IP and TCP/UDP. The high level tests are those that look at higher layer protocols such as LPR. In the following, we begin with low level tests and work upward.

5.2.1 Low level attacks

Six well-known low level attacks based on vulnerabilities in IP were run as tests against the DC 230ST. They were all denial of service attacks that had worked against other operating systems in the past. The attacks are listed below.

- bonk
- jolt
- nestea
- newtear
- syndrop
- teardrop

All of these attacks failed to crash the DC 230ST, but we did lose IP network access to the DC 230ST during the jolt attack. This is most likely caused by finite IP fragment re-assembly resources on the Docucenter. The jolt attack sends many IP fragments to the target and may be consuming a fixed resource. This attack could be used to deny access to the Docucenter, but there are other ways to do the same thing (e.g. TCP SYN floods).

We also tested the DC230ST with three slightly higher level attacks called land, winnuke, and ping flood. The only noticeable effect on the Docucenter was the normal network slowdown caused by the ping flood.

An interesting feature of the Docucenter's low level network implementation was discovered while performing the port scans described in Section 5.2.2. The DC 230ST's IP implementation has trivially guessable sequence numbers. Although conformant to the specification for IP, this greatly simplifies the spoofing of TCP connections destined for the Docucenter. If the DC 230ST had IP-based access control, this would allow an attacker to spoof an authorized host's IP address in a TCP connection, possibly allowing unauthorized submission of print jobs, etc. As will be discussed below, the DC 230ST does not have IP-based access control but the issue of guessable sequence numbers is still a serious concern.

To summarize the results of our low level network testing, no serious compromises were found. In our opinion, the IP protocol stack implementation in the DC 230ST is either well maintained, and therefore patched to eliminate these vulnerabilities, or well implemented from the start. With the exception of guessable sequence numbers, we were quite satisfied with its IP protocol stack.

5.2.2 Port scan

The `nmap` [Fyo] tool was used to scan for open TCP and UDP services. The results were as follows:

TCP		
Port	service name	description
25	smtp	Electronic Mail
79	finger	List current interactive users
80	httpd	Web server
111	sunrpc	Remote Procedure Calls
514	shell	Remote Shell
515	printer	Network Printer Daemon
1024	unknown	

1. **smtp** - This service is used when a machine is expected to *receive* electronic mail. None of the documentation for the DC 230ST explains why this port is open. Many of the common daemons that run on this port have been plagued with remote attacks. Connections to this port from our local subnet were closed immediately. We hypothesize that this port may only accept connections from a particular set of IP addresses (e.g. Xerox technical support sites).
2. **finger** - This service is only useful on machines that have interactive users. The DC 230ST documentation does not explain why this port is open. Many of the common daemons that run on this port have been plagued with remote attacks. Connections to this port from our local subnet were closed immediately. We hypothesize that this port may only accept connections from a particular set of IP addresses (e.g. Xerox technical support sites).
3. **httpd** - This service is used by a web server to handle requests for web pages. The DC 230ST uses the web server to allow remote users to upload files (text, postscript, or PCL) and request that they be printed. There is no authentication of users, nor restriction of IP addresses. Anyone that can connect to this port can upload and print files. All transactions are charged to a generic *webuser* account.
4. **sunrpc** - This service is used to remotely run applications from "trusted" hosts without authentication. The DC 230ST documentation does not explain why this port is open.
5. **shell** - This service is used to establish a remote interactive session, usually from a "trusted" host. The DC 230ST documentation does not explain why this port is open. Connections to this port from our local subnet were closed immediately. We hypothesize that this port may only accept connections from a particular set of IP addresses (e.g. Xerox technical support sites).
6. **printer** - This service is used to establish the network printer functionality of the DC 230ST for systems that use the LPR printing protocol. There is no authentication of users, nor restriction of IP addresses. Anyone that can connect to this port can upload and print files.

UDP		
Port	service name	description
111	sunrpc	Remote Procedure Calls
161	snmp	Simple Network Management Protocol
518	ntalk	Network talk (chat program)
1026	unknown	Unknown RPC service
1029	unknown	Unknown RPC service

1. **sunrpc** - Used to request remote services from "trusted" hosts without authentication. The DC 230ST documentation does not explain why this port is open.
2. **snmp** - A protocol used to remotely administrate networked components.
3. **ntalk** - This port is usually used for some types of network *chat* programs. The DC 230ST documentation does not explain why this port is open.

5.2.3 ISS Scan

We scanned the DC 230ST using the ISS security scanning tool. All known attacks were selected to be scanned. ISS detected nothing other than the existence of the previously mentioned services that expose the machine to risk of attack.

5.2.4 SNMP

We found that an attacker can query the DC 230ST to collect information via SNMP . This information includes nearly all of the information available from the system's console panel including network settings, the Banyan Vines user names and passwords, and the system's physical location. An attacker could use this information for password guessing on other systems, to discover other network information, etc.

To make matters worse, some of these setting can be changed using SNMP because the machine has the default private community string of "private". Using the Linux port of the Carnegie Mellon University SNMP package [SS$^+$], we were able to change the settings of several system values including the Banyan Vines user name and password! Other changeable values included the physical location field and the TCP service port of the LPR service. Attempts to change other network parameters such as the IP address or network mask in this manner failed. But this does not imply that a more sophisticated attack would fail. Finally, we could find no method for changing the private community string in the documentation.

One implication of the SNMP facility of the DC 230ST is that the system must be protected by firewall mechanisms to prevent unauthorized access to the system. Furthermore, the impact of SNMP access to the DC 230ST may be much worse than we have presented as there are over 600 SNMP entries available in the Docucenter specific variables, and we could readily identify the meaning of fewer than 50.

5.2.5 HTTP

By sniffing the exchange that takes place during a print request via the web server, we were able to determine that the HTTP server is a **MicroServer** by Spyglass, Inc [Spy]. This product is designed to be a low footprint (10–36 KB needed) web server used in hardware devices such as copiers and fax machines.

5.2.6 LPR

There are a number of attacks that can be made via the network printer service. We tested a few of the common buffer overflow attacks with no success. As part of the specification for this protocol [III90], a remote user is supposed to be able to request that arbitrary files on the LPR server be printed or deleted. Our attempts at exploiting this feature of the protocol were unsuccessful.

5.2.7 Lack of Access Control

The most discouraging problem with the DC 230ST's IP-based high level protocols is the lack of access control. These protocols, such as LPR and HTTP, can be disabled, but there is no built-in mechanism for preventing unauthorized users, with access to the network, from submitting print jobs, checking the status of print jobs, etc. One possible mechanism would be to allow the administrator to specify certain groups of IP addresses as authorized to access a given service. While this approach has some problems, it is better than nothing and requires no changes to existing network protocols. Indeed, in an Internet-connected or WAN-connected environment, the DC 230ST should be put behind a firewall mechanism that is configured to block outside access to all ports on the DC 230 ST.

5.3 Physical

5.3.1 Control Panel Access

The DC 230ST has a touch-sensitive control panel and a keypad on the top of the unit. This is the typical interface used to make copies, fax paper documents, and initially configure the system. Access to the documented configuration menus is restricted using a numeric password which is "22222" by default. Interestingly, the value for this password given in the documentation is incorrect. We found the password by calling Xerox technical support. This password can and should be changed from the control panel. The procedure for using the password is to push the "Access" key, enter the password, and press the return button.

A more interesting feature of the DC 230ST is an undocumented maintenance system which is accessible from the control panel. The procedure for accessing the maintenance system is the same as for the configuration menus, but the password is "#11". This password

does not appear to be changeable and therefore is an open backdoor for those with physical access.

The maintenance system allows the user to modify the non-volatile random access memory (NVRAM) in the system, run hardware and software diagnostic programs, and change settings for some system features. For instance, we successfully used this system to turn the power saving mode of the DC 230ST off. One interesting feature is the ability to modify NVRAM in a straightforward manner. This could be used to modify the boot up sequence and its parameters thereby compromising the Docucenter.

5.3.2 Serial Access

We tested both serial interfaces discussed in Section 3. The one underneath the control panel was unresponsive, but we determined, through social engineering, that this port was used for updating the software that manages the control panel.

The second serial interface was much more interesting. We connected a terminal to the port and were immediately greeted with a login prompt from the control computer. Many attempts at guessing a user name-password pair were tried unsuccessfully. However, we are relatively certain that all DC 230ST's have the same passwords loaded on them. We base this on a conversation we had with a Xerox technical support person. We were informed that the internal hard drive could be ordered preloaded with all software and ready for installation. This is so a corrupted drive can be easily replaced. Further, efforts could be made to build an automated password guesser to try to break these passwords or to read the hard drive and obtain the password hashes that could then be fed to software such as crack.

Upon restarting the system, we found that it had an NVRAM setup phase similar to the CMOS setup on most PC's. The NVRAM had three access options "Novice," "Administrator," and "Advanced." Novice had no password, but provided only limited functionality. The other two accounts were poorly passworded. The password account for Advanced was "Advanced" and similarly for the Administrator account. These modes allow the user to execute arbitrary boot loader code in a language similar to FORTH. The accounts also allowed the user to modify NVRAM.

The serial interface allows anyone to rewrite the entire operating system on the hard disk. We closely observed this process being done by Xerox support. During this process, command messages are sent via the serial port and the bulk data flow is through the parallel port.

The serial port interface provides a strong mechanism for someone with physical access to implant a trojan horse into the DC 230ST. The result could be nearly arbitrary cross-channel communication and monitoring.

5.3.3 Floppy Disk Access

A 3.5 inch floppy disk drive is integrated into the DC 230ST for two purposes. The first purpose is to allow users to print documents from disk. This is done by inserting the disk with the print file on it and selecting the file to print using the control panel. The second purpose is to allow an administrator to add functionality to the DC 230ST. In this case, an upgrade disk is provided by Xerox which is inserted into the drive. After entering the administrative password into the control panel, one option is to upgrade from disk. Even the seemingly innocent print from disk feature entails some risk. Anyone with physical access can insert a disk and print a file that may contain postscript viruses. These viruses can then store documents printed or faxed for later retrieval.

Assuming that the administrative password is changed and kept secret, the disk upgrade mechanism is of less concern.

5.3.4 Hardware Access

Physical access to the DC 230ST must be carefully controlled. A good example of this is the vulnerability of a hard drive swap. The hard drive on our DC 230ST was a 2 Gigabyte SCSI Quantum Fireball II. A well equipped attacker could create a hard drive with embedded monitoring or relaying software and swap it during somewhat extended access to the machine. The development of such a trojaned system disk would not be difficult given sufficient resources.

6 Conclusions and recommendations

In this section we present our conclusions about the DC 230ST and our recommendations for addressing the concerns raised in those conclusions.

The DC 230ST has a large number of documented and undocumented network services, but there is no way to prevent unauthorized users or hosts from using most of these services. In the case of services that allow submission of print jobs, these services would allow any individual with network access to submit a Postscript virus that may monitor or relay network print jobs and faxes. We recommend that the DC 230ST be contained within a packet filtering firewall of its own, thereby preventing unauthorized hosts from accessing the Docucenter. The firewall mechanism should allow incoming connections from authorized hosts for the following services (if needed) : LPR (TCP port 515), HTTP (TCP port 80), and NetBIOS (TCP port 139) Other ports may need opened for specific services not mentioned here, but these should be opened only upon careful consideration.

The integrated hard drive in the DC 230ST allows unauthorized reprogramming of the Docucenter, assuming physical access. It should be noted that known mechanisms exist, such as the serial port interface, by which the system can be easily reprogrammed without extreme physical access such as disassembly. It is also possible to build a replacement hard drive and substitute it given internal access to the DC 230ST. We recommend that strong

measures be taken to assure secure physical access. Also, we recommend changing all default NVRAM and control panel passwords to further thwart physical access attacks.

Fortunately, the FAX and network services of the DC 230ST appear resistant to many types of attack. The Docucenter's operating system appears to be well implemented and fairly robust.

Because our preliminary attempts to access the Docucenter's operating system via the serial port failed, we have been unable to do any significant white box testing. We recommend future work on the DC 230ST, involving the ability to log onto the system. The passwords required might be guessed by a brute force password guesser. This would allow more detailed and in-depth testing of the system. It might also allow discovery of the nature of the undocumented network services.

During our use of the DC 230ST, we noticed a definite problem with general stability. Occasionally, the internal computer would crash leaving the network printing and FAX services off-line. This was nearly always solved by restarting the system. We were never able to discern the exact problem. Some crashes seemed to coincide with the automatic power down features of the system or long periods of system inactivity.

Finally, the supplied documentation is incorrect in some cases. We experienced incorrect default passwords and incorrect installation instructions for accompanying software. We recommend that an effort is made by the deployer of this technology to keep a set of corrections to the documentation that accompanies the DC 230ST.

In conclusion, the DC 230ST appears resistant to many high-level network attacks, but it has little concept of access control for users and hosts. Use in a highly sensitive environment mandates strict physical security to prevent system compromise.

References

[Doe] Gert Doering. mgetty+sendfax. `http://www.leo.org/~doering/mgetty/index.html`.

[Fyo] Fyodor. nmap – the network mapper. Available via HTTP. `http://www.insecure.org/nmap/`.

[III90] L. McLaughlin III. *RFC 1179 - Line Printer Daemon Protocol*. The Wollongong Group, August 1990.

[Lin75] Richard Linde. Operating system penetration. In *National Computer Conference*, pages 361–368, 1975.

[OL95] William J. Orvis and Allan L. Van Lehn. Data security vulnerabilities of facsimile machines and digital copiers. Technical Report UCRL-AR-118607/CIAC 2304, Department of Energy, Computer Incident Advisory Capability (CIAC), Lawrence Livermore National Laboratory, January 1995.

[Spy] Spyglass, Inc. *Spyglass MicroServer 2.0.* `http://www.spyglass.com/solutions/technologies/microserver/`.

[SS+] Jürgen Schönwälder, Erik Schönfelder, et al. Linux cmu snmp project. `http://www.gaertner.de/snmp/`.

[TDS99] Meeting with Trident Data Systems. Personal communication, May 13 1999.