Special Issue on Printers: UDC 681.612.5.064-185.4

Internet Printing

Philip THAMBIDURAI*, Tim DEPPA*, Kazuo NAKANO**

Abstract

The rapid growth of the Internet will significantly impact the requirements for future printing devices. The ability to communicate with anyone on the Internet at low cost provides new channels for transmitting and printing high quality documents. In this paper we discuss "Internet Printing" and focus on an industry-wide effort at developing a common protocol.

1. Introduction

The internet is a highly reliable media which allows the transmission, spread and exchange of information quickly at low cost. The Internet has made it easy for users to access a tremendous store of information on virtually any topic. This base of information has led to a large increase in the volume of printed paper, because a user prints not only those documents that he/she created but also other documents that include Web pages. As high speed access to the Internet becomes the norm, the Internet will become as ubiquitous and robust as the telephone system, and even more documents will be printed in the near future.

This is because paper is the desired medium for reading and consuming information, although the information may be seen on a monitor.

Conventionally business documents have been first printed at the site where the document is created and then sent via the postal services to the destinations. If dissemination must occur instantly, facsimile transmission is the norm, although at the expense of document quality and high telephone charges. If time is not critical, documents can be sent via e-mail, however most companies impose data length limitations on e-mail documents, so it is the norm that a high quality document cannot be attached.

A tantalizing possibility is to send documents directly to printers worldwide, without sacrificing document quality & color, time, or cost.

Printing in the enterprise or corporate environment today involves a multitude of protocols. One of the primary driving factors behind an "Internet Printing" protocol is to allow companies to standardize on a single protocol which takes advantage of the emerging Internet/Intranet infrastructure. It would be possible to print out documents to a printer that is located in another domain or company under a consistent protocol, and even across enterprises via the Internet.

Internet Printing can thus be viewed as an extension of enterprise printing

Figure 1 shows the concept of Internet printing.

* System Engineering Group, Oki Data Corporation
** Vice President, Engineering Division, Oki Data Corporation

2. Enterprise Printing Architectures

A variety of printing technologies have evolved to meet the needs of the enterprise, which encompasses both LANs (Local Area Networks) and WANs (Wide Area Networks). Table 1 summarizes the basic methods for attaching printers to LANs, and the printing methods, protocols, and operating systems.

Network printing has grown in complexity in part due to the unique requirements which each Network Operating System (NOS) imposes, and in part due to the proprietary technologies offered by each printer vendor for management and configuration.

For example, Novell and Microsoft[*4] each provide their own directory and print services technologies, and printer vendors use a variety of managment protocols to communicate with their internal and external (Network Interface Card) NICs. Typical management protocols include the Network Printing Alliance Protocol (NPAP), Network Printer Task Force (NPTF) protocol, Data Link Control (DLC) protocol, and the Novell Service Access Protocol (NSAP).

Each printer vendor must therefore bundle its own proprietary software to configure and manage network printers. Given this complexity, it is not surprising that network administrators continue to identify network printer management as their most difficult task.

The Internet has created new requirements for management, directory services, printing protocols, and security which extend beyond the enterprise. It has also provided the

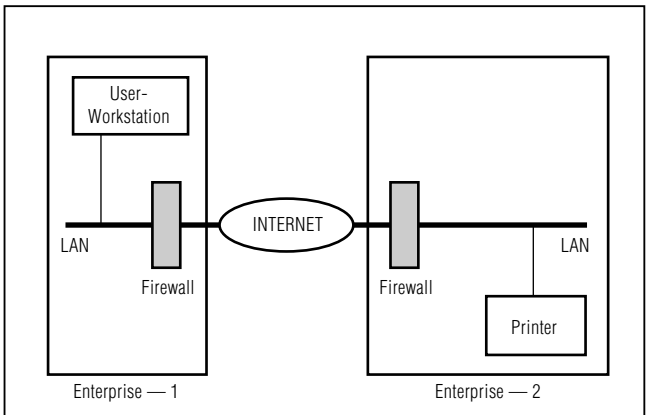

Figure 1: Internet Printing Concept

impetus for standardizing network printer management software and protocols. Although the printer industry had finalized on a standard for SNMP (Simple Network Management Protocol) management of printers, vendors continued to sell proprietary management software. However, the popularity of the Internet has recently influenced printer vendors to offer Web browser-based management utilities, which include SNMP support.

The browser-based approach makes it almost automatic to support most operating systems. In addition, since SNMP is itself a basic Internet standard, it has become a mandatory requirement for network printers.

The combination of browser-based management and SNMP has finally made it possible to manage all enterprise printers in a uniform way. In effect, the Internet is becoming the driving force for the future direction of enterprise printing.

3. Unique Requirements for Internet Printing

Ideally an Internet printing protocol[1,2,3] should take advantage of existing Internet protocols and standards. Using the existing Web infrastructure will lead to shorter implementation times and allow the printing protocol to focus on the printing environment instead of lower-level transport protocols and issues. Among the most widely used protocols in the Internet today are TCP/IP (Transmission Control Protocol / Internet Protocol) and HTTP (Hyper Text Transfer Protocol). In a strict sense, TCP is a transport level protocol and HTTP is an application level protocol.

From an end-user perspective, the requirements for Internet Printing are for the most part similar to those for enterprise printing. End-users requirements include viewing printer and job status, job transfer and job request, and obtaining a list of pending jobs.

System Administrators must be able to monitor, configure, and control the printer from remote sites. End users must have the same conveniences as printing to a locally attached printer or a network printer.

While printing to a network printer within the Intranet, security is typically not an issue. The printer is relatively safe from attempts to "break-in" from outside the company, and documents are not encrypted before transmission to the printer. Allowing access to the printer (or print server) from anyone in the Internet requires both authentication and encryption, in general.

Almost every company today that connects to the Internet uses a firewall to allow only certain protocols, usually based on TCP or UDP (User Datagram Protocol) port numbers. An entirely new protocol will have a new port number and firewall administrators must explicitly configure their firewalls to allow traffic on this port.

Another requirement commonly not found in typical network printing environments is a directory service for locating printers and identifying their capabilities. A directory service will allow an end-user to determine which printers are available. The directory schema could also describe printer properties. This service is similar to a "yellow pages" for looking up telephone numbers. While there is much ongoing activity in Internet directories including existing standards, the adoption of these standards across the enterprise has been slow, partly due to the complexity involved in setting up the service.

Querying a printer that resides in another country could possibly include the need to handle foreign languages in some graceful manner, as far as human-readable text strings are concerned. A printer in a foreign domain might provide its human-readable status in some local language. The user or client operating system must be informed of the character set and other language attributes. This will also allow the operating system to suppress these strings if the language is not supported. It is thus necessary for any Internet protocol that interacts with the user via text strings or names to consider Internationalization issues (languages, character sets, etc.).

The typical method of printing from the MS-Windows environment requires the local printer driver to be installed in the user's computer. The driver generates the PDL (Page Description Language) data stream, which the printer can understand; this is most often the PCL (Printer Control Language)[*7] format, but could also be PostScript or other.

When printing to a remote Internet printer it will still be necessary for a driver to generate the PDL. Whether the printer driver should be installed locally in the end-user's host or in the printer site, is really left to the implementation. In some cases, a local printer driver does not have to be installed. For instance, in a server-based implementation of IPP (the Internet Printing Protocol), it is possible to have the driver reside in the server and the user's document could be transmitted in "Windows enhanced metafile format" this will make it unnecessary to install a local print driver.

Printer Attachment to LAN	Printing Methods	Protocols	Network Operating Systems	Client Operating Systems
• Client workstation parallel port	• Print Server / Spooler	• Novell IPX/SPX	• Novell Netware	• Win95
• File Server parallel port	• Peer-to-Peer	• TCP/IP *2)	• Win NT	• Win 3.1x
• Internal Network Interface Card (NIC) • External NIC		• EtherTalk *3) • Win NT *4)	• UNIX *5) • Mac	• Mac• UNIX
		NetBEUI/DLC	• OS/2 *6)	
			• Windows *4) for Workgroups	

Table 1: Enterprise Printing

*1. Novell IPX/5PX is a registered trademark of Novel Co. *2. TCP/IP is a protocol developed by the US Defence Agency. *3. EtherTalk, Mac are registered trademarks of Apple Computer Inc. *4. Microsoft, Windows NT and Windows are trademarks or registered trademarks of Microsoft Corporation in the USA and other countries.b *5. UNIX is a registered trademark licensed by X/Open Company Ltd. in the USA and other countries. *6. OS/2 is a registered trademark of IBM Corporation. *7. PCL is a registered Trademark of Hewlett-Packard Co.

4. Approaches to Internet Printing

Early attempts to implement minimal Internet Printing capability have been based on using the Simple Mail Transfer Protocol (SMTP, the e-mail transport protocol), where the document or image to be printed is sent as a MIME (Multipurpose Internet Mail Exchange) attachment to an e-mail message. A significant drawback of this approach is that the user cannot interact with the printer, and does not know when the e-mail message will be delivered to the printer. Essentially, the user cannot interact with the printer. Besides, the e-mail infrastructure as it exists today is not sufficiently robust. Another disadvantage is that implementations by different vendors even though based on SMTP are not interoperable because of non-standard vendor-specific codes sent in the text portion of the e-mail message.

Some vendors have defined their own protocols for Internet Printing based on TCP. These are proprietary and no attempt has been made at standardization. Currently even if both the sending and receiving ends support this protocol, the firewalls typically will not allow the protocol to pass through.

In the Fall of 1997 Hewlett-Packard[*8] introduced the JetSend[*8] protocol for transmission of images/data between simple devices[4]. The JetSend specification itself is independent of the transport protocol. JetSend is intended for sending images from one "simple" device to another, based on the data-handling capabilities of the receiving device.

JetSend will allow a digital camera to send an image directly to the printer without first going through a personal computer. Devices that are JetSend-enabled will still require some device specific code. At the present time, the JetSend protocol does not explicitly address issues related to the Internet such as security and naming. It is currently best suited for operation on a local area network or direct connection such as IEEE 1394. However, JetSend technology has the potential to support Internet Printing in the future. For now, IPP and JetSend address quite different requirements.

5. The Internet Printing Protocol (IPP)

IPP is in the final stages of development under a working group of the IETF (the Internet Engineering Task Force, which is one of the standards bodies for Internet protocols). The working group includes representatives from major operating system and printer vendors.

IPP is basically a simple client/server protocol. The client is the end-user application, while the server is the printer or a print-server, which may be implemented in software on a workstation. Printers may directly embed IPP support or they might communicate with a server using a proprietary protocol, in which case the server is the end-point or "IPP printer" to the end-user.

The first version of IPP focuses on end-user requirements, leaving administrative and management issues for the next version. End-user operations include:

The most common end-user operations would be to print a document, cancel a job, and to view the job status. The Validate operation is used to test if the printer and document attributes desired by the user for the particular document are fully supported by the destination printer. It simulates a print job but does not actually send any data to the printer; the printer would respond by informing the user if the desired attributes are supported. When the job completes the printer can notify the user.

Figure 2 shows IPP configurations. Using IPP, an end-user can print in three ways:

1. From within an application. In the MS-Windows environment, this is the most common method of printing. A user typically selects the "print" option and is then provided with the default and other printers that have been installed.
2. Sending a pre-formatted file to the printer. The file might be in PostScript[*9] format or other page description language which the printer can understand.
3. Providing a reference to the document (a URL — Uniform Resource Locator). The printer will retrieve the document from the Internet site specified in the URL, either by using HTTP or FTP. This operation is optional in Version 1.

An IPP printer is identified using a URL. This allows for a uniform naming scheme that is consistent with Web addresses. A large number of printer and job attributes are defined in IPP. These attributes allow the user or operator to specify exactly the desired properties of a print job. As much as possible, these attributes are aligned with those defined in the Printer MIB (Printer Management Information Base) and the Job MIB.

In terms of security, IPP specifies the use of Transport Level Security (TLS) which is an emerging standard, used for both encryption and authentication. Note than authentication could be required for both the user and the printer. The first version of TLS is compatible

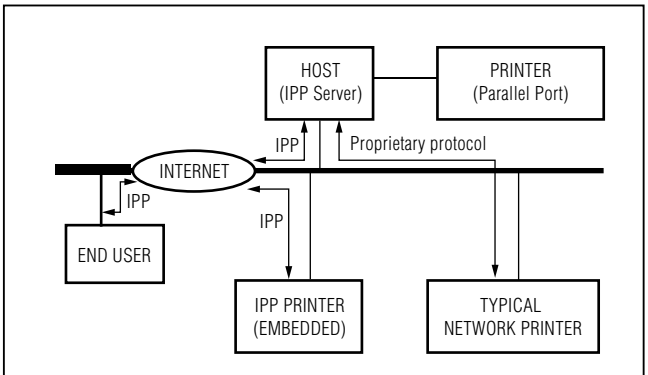

Figure 2: IPP Configurations

*8. Hewlett-Packard, JetSend are trademarks or registered trademarks of Hewlett-Packard Co.

*9. PostScript is a registered trademark of Adobe Systems Inc.

with the Secure Sockets Layer (SSL 3.0) standard, which is the most commonly used security scheme in HTTP today. TLS has recently been submitted to the IETF (Internet Engineering Task Force) as a proposed standard. It is expected to be the mainstay for HTTP security. The basic intent of the IPP working group is to use the security protocols being defined for HTTP, instead of inventing a new one. A particular implementation of IPP could offer various levels of security, including none at all. The appropriate security level is negotiated while establishing the HTTP session.

Implementing IPP is not difficult since it uses simple HTTP posts. At the end-user side HTTP is available in all modern operating systems via the browser and via language APIs such as Java*10, the C/C++ Win-32 Internet API. At the server end, the HTTP server that must be implemented is very basic, unlike typical Web servers. Most of the complexity of implementing HTTP servers resides in the "proxy servers", which is the component needed for supporting IPP, and not the "origin servers". Embedded Web servers are available today with a memory footprint of less than one Kbyte. It is likely that many vendors will embed IPP support in the printer or a hardware print server. Other vendors will not do anything at all — a software-based print server running on a workstation or host machine will serve as the end point for IPP; the printer could communicate with this IPP server using existing proprietary protocols without any change.

IPP is expected to become an IETF standard in 1998. Support for IPP in operating systems and printers is expected in late 1998. Currently several vendors are implementing and testing IPP prototypes, both embedded and server-based. Testing interoperability between different vendor implementations has also begun, albeit based on the yet-to-be standard version of IPP. Early indications are that Microsoft will support IPP in future implementations of the Windows architecture.

6. Conclusion

A common standard for Internet Printing and Intranet Printing provides significant benefits to end-users. Apart from being a common method of printing from diverse operating systems and platforms, it also takes advantage of the Web infrastructure to extend the concept of network printing into the Internet. If an Internet Printing protocol is widely implemented and supported, it could also serve as a low-cost method for transmitting high-quality documents, possibly rivaling the existing telephone line based FAX machines in the enterprise. The Internet will cause enterprises to focus on standards and interoperability, instead of a proprietary protocol.

As the Internet evolves, there are still many difficult issues and infrastructure problems that need to be addressed. These include enterprise management, directory services, document formats, and security. It is likely that these and other issues will be solved first for the Internet, and then become the standards for enterprise printing.

Okidata is currently following the developments in this area, especially the IPP and JetSend protocols. These will be supported in future network products if the markets demand it, and printers will be much more than simple output devices. We must consider how to increase the added value of printers and how to deliver benefits to end users.

7. References

1. IPP Model and Semantics, Version 1.0, Internet Draft document, October 27, 1997.
2. IPP Protocol Specification, Version 1.0, Internet Draft document, October 23, 1997.
3. IPP Requirements, Version 1.0, Internet Draft document, October 15, 1997.
4. Peripherals Win A Voice of Their Own with HP's JetSend, The Hard Copy Observer Page 1, (August 1997): 34-37.

*10. Java is a registered trademark of Sun Microsystems Co.

*. Other product names and company names are registered trademarks or trademars of respective companies.