



UNDERSTANDING AND MEASURING EXPECTED PERFORMANCE: UNIPRINT SERVER VERSION 5.0 AND UNIPRINT GATEWAY VERSION 5.0

TECHNICAL WHITE PAPER



Copyright and Trademarks

Copyright ©2005 INGENICA. a division of Bell Business Solutions This documentation cannot be reproduced in full or in part by any means without our prior written consent.

We provide this documentation as is without warranty of any kind, express or implied, including but not limited to implied warranties of appropriateness for specific purposes or merchantability. In no event or under any circumstances shall we or our suppliers or distributors be liable for any damages whatsoever, including and without limitation, damages resulting from business loss, which may arise from the use or inability to use this documentation, even if we, our suppliers or our distributors have been previously advised of the possibility of such damages. Since some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

Our name and logo are trademarks. Other brand and product names are the trademarks or the registered trademarks of their respective corporations.

Table of Contents

| | |
|---|------------|
| Abstract | iii |
| Introduction | 1 |
| UniPrint Dataflow Review | 1 |
| Measurable Time Points | 1 |
| How to Measure the Expected Times of Each Activity | 2 |
| Job Spooling Time | 3 |
| UniPrint Server | 3 |
| UniPrint Gateway | 3 |
| Job Polling Cycle | 4 |
| UniPrint Server | 4 |
| UniPrint Client | 4 |
| UniPrint Gateway | 5 |
| UniPrint Print Server | 5 |
| Job Transfer Time | 6 |
| UniPrint Server | 8 |
| UniPrint Gateway | 8 |

Abstract

With all UniPrint components functioning as designed, the number one cause of lengthy overall print speed is Job Transfer Time affected by bandwidth contention. Network bottlenecks usually occur intermittently – slow during periods of high contention, but suddenly disappear during periods of low contention. Optimizing print job file size helps to reduce contention but cannot eliminate contention caused by other activities competing for the LAN, WAN, or Internet. UniPrint processes that are not network related such as local Job Spooling and local Job Polling have clearly defined performance expectations and are easily verified, all things being equal.

The significant differences between UniPrint and Windows printing are file size and transfer protocol. UniPrint Server produces PDF files which are significantly smaller than PCL, Postscript, or EMF files. UniPrint Gateway produces files identical in size to UniPrint Server but achieves better network throughput. UniPrint Gateway uses FTP to move print jobs compared with traditional network printing which uses Microsoft's Server Message Block (SMB) protocol. FTP is faster (achieves better throughput) than SMB such as file size and network conditions.

Introduction

The following discussion provides information useful for troubleshooting overall UniPrint print times and for verifying expected UniPrint performance from print to printer. The information is organized according to the three main time-based activities performed by the UniPrint processes: Job Spooling, Job Polling, and Job Transfer. Each activity is examined in both a UniPrint Server and UniPrint Gateway environment. Instructions are also provided for measuring expected performance of each activity as they relate to the various UniPrint components.

UniPrint Dataflow Review

The data flow of a UniPrint print job varies depending on the UniPrint product being used. Printing to the “UniPrint” printer within a Server-based Computing (SbC) session, creates the print job on the host server and returns it directly to the connected client’s workstation encapsulated within the thin-protocol being used, ICA or RDP. If UniPrint Gateway is being used, then users print to familiar named instances of the UniPrint driver also known as *Gateway Printers*; for example, “PrtSrv - HP Laser Jet 5M”. Gateway printers create print jobs in exactly the same way as the *UniPrint* printer but send them (using FTP) to a centralized Spool Server. Gateway Print Servers (Windows 2000/2003/XP with UniPrint Print Server software installed) retrieve print jobs from the Spool Server using FTP.

Measurable Time Points

There are three distinct activities that can be measured against time for all UniPrint jobs:

1. Job Spooling Time
2. Job Polling Cycle
3. Job Transfer Time

Job Spooling Time is the time it takes the UniPrint driver to create the print job and write it to the host server’s spool location. The Job Polling Cycle is a period of time between which certain UniPrint processes check for spooled jobs and begin to transfer them off the host server to the appropriate network path. In the case of UniPrint Server, the job polling cycle is non-configurable and preset to 2000 ms. The job polling rate for UniPrint Gateway is configurable through the UniPrint Configurator and defaults to ten seconds¹. Job Transfer Time is how long a UniPrint print job takes to move from one UniPrint component to the next over the LAN, WAN, or Internet. The table below identifies each UniPrint component that performs one or all of the above activities.

¹ The job polling cycle works in conjunction with a configurable buffer threshold which is set to 30K by default. If the buffer threshold is $\geq 30K$, then a job transfer is initiated immediately. Both settings can be configured through the UniPrint Configurator. See the Job Polling Cycle section below for more details.

| Job Spooling Time | Job Polling Cycle | Job Transfer Time |
|-------------------------|-----------------------|---|
| UniPrint Server | | |
| Terminal Server | Terminal Server | Terminal Server to Client Workstation |
| Client Workstation | Client Workstation | Client Workstation to LAN Print Server ^a |
| UniPrint Gateway | | |
| Terminal Server | Terminal Server | Terminal Server to Spool Server |
| UniPrint Print Server | UniPrint Print Server | Spool Server to UniPrint Print Server |

a. The UniPrint Client processes jobs using locally installed print drivers. If the destination print device is a network printer, then regular LAN printing, e.g., SMB printing is used from this point onward.

The overall time for a print job to appear as a *hard copy* printout may vary depending on how well each of the above activities performs. Each activity's performance is affected by a variety of factors. The UniPrint driver must contend with server-based sessions for CPU, RAM and hard disk resources. UniPrint Gateway must contend with other network traffic when transferring jobs from the UniPrint Spool Server to the UniPrint Print Server. The application from which a user prints can affect the job spool time. For example, a simple Word 2000 document with a variety of fonts, formatting, and embedded graphics typically will spool faster using UniPrint than a native print driver processing the same job. However, a large PowerPoint 2000 presentation with complicated graphics (background images, multiple gradients, etc.) may spool faster using a native PCL print driver. While this is not true for all documents of this type, applications do play a significant role in Windows printing. Applications developed prior to various Windows printing subsystem enhancements are likely to use legacy subsystem functions and thus experience reduced spooling performance. These applications affect all print drivers, UniPrint and native drivers. The same PowerPoint presentation printed from PowerPoint 2003 will print faster than its predecessor creating a smaller print job. In this case, UniPrint matches or outperforms native print drivers when spooling jobs printed from the same application.

In most cases, optimizing server-based printing involves identifying, prioritizing, and managing a finite set of resources, CPU, RAM, Network, Disk I/O. Understanding how UniPrint uses these resources to create and route print jobs to destination printers will allow administrators to measure the impact of UniPrint compared to other printing solutions.

How to Measure the Expected Times of Each Activity

To ensure accurate measurements, keep the following constant for each set of tests:

1. Use that same application
2. Use the same document

3. Use the same printer (or print server)

Only vary the above once you have established baseline measurements that represent typical printing scenarios within your Server-based Computing environment. For example, if your organization uses Office 2003, do not conduct your tests with Office 2000 applications. If you suspect that Job Spooling is not performing as expected, then vary the document or application version and compare the results with your baseline measurements.

Job Spooling Time

Job Spooling Time is the time it takes the UniPrint Driver to create the print job and write it to the host server's spool location. Job spooling time occurs on the Terminal Server, client workstation and UniPrint Print Server, if using UniPrint Gateway.

UniPrint Server

On the host Citrix server, UniPrint spools its jobs to a directory. The directory created depends on the File Name settings chosen by the administrator. These settings are set through the Spooling tab in the UniPrint Configurator. If user name and session ID are both selected, then UniPrint spools its jobs to the following directory: %systemroot%\system32\spool\PRINTERS\UniPrint\%username% \<sessionID>, for example, "C:\Windows\system32\spool\PRINTERS\UniPrint\jdoe\3". The session ID will be the current connected user's session ID as assigned by Terminal Services.

1. Connect to a desktop session as Administrator and open a window to the above path so that you can see the print job spool to this folder when you print.
2. Print to *UniPrint* and measure the time it takes for the job to fully spool to the above location. Two files should be spooled to this directory: the actual print job *.ppf, and a *.pcf file. The PPF file is the actual PDF file, and the PCF file is a text file containing the UniPrint printer property selections among other information. This file should only be 2K

Measure the time it takes to fully spool the job on the Terminal Server. Use a large enough file so that the measurement can be recorded in seconds using a watch. This time period can vary depending on the application being used and the size and type of the document being printed. Set your window view to "Detail" so that you can see the reported size of the spooled file. You may wish to repeat this procedure several times varying the UniPrint driver selections until you have determined the optimal settings for the type of document being printed. Consult the UniPrint Administrators Guide for explanations of the configurable settings.

Once you have optimized the UniPrint driver settings, establish a baseline by measuring the time between clicking on the Print button to when the job is fully spooled to the directory location above (end the measurement just before the file is transferred into the virtual channel). UniPrint uses the same Windows printing subsystem components used by native drivers. To compare the job spooling efficiency of UniPrint against other printing methods, e.g., auto-created printers, perform the same measurement using an auto-created printer, the same application, similar driver settings, and the same document.

UniPrint Gateway

On the host Citrix server, *Gateway Printers* spool their jobs to the following directory:

%systemroot%\system32\spool\PRINTERS\UniPrint\

1. Connect to a desktop session as Administrator and open a window to the above path so that you can see the print job spool to this folder when you print.
2. Print to a Gateway printer prefixed with <Print_Server_Hostname> (for example, PrtSrv001 – HP LaserJet 5M) and measure the time it takes for the job to fully spool to the above location. Two files will be spooled to this directory: the actual print job *.ppf, and a *.pcf file. The PPF file is the actual PDF file, and the PCF file is a text file containing the UniPrint printer property selections among other information. This file should only be 2K.
3. If you wish to trap the print job (prevent it from being transferred off of the Terminal Server), then on the Terminal Server, open Services in Control Panel and stop the “UniPrint Server Plus – Upload Manager” service.

All things being equal², the time measurement here should match exactly that of the UniPrint printer as outlined in the previous test above. Gateway printers are instances of the same UniPrint driver and therefore should spool jobs identically.

Job Polling Cycle

The Job Polling Cycle determines the interval time between checks for spooled UniPrint jobs. Job polling occurs on the Terminal Server, Print Server (using UniPrint Gateway) and UniPrint Client (using UniPrint Server) components.

UniPrint Server

The rate at which UniPrint Server checks for print jobs in the %systemroot%\system32\spool\PRINTERS\UniPrint\%username%\<sessionID> folder is every 2000ms. To test this time period and to determine if print job transfers are being initiated in a timely fashion, repeat the Job Spooling Time experiment for UniPrint Server as outlined above. After the print job has been completely spooled to the above folder, the process responsible for initiating the network transfer should occur rapidly (within 2000ms). This time period can be verified by measuring the time from when the job is completely spooled to the appearance of the virtual channel spool progress bar (provided it has not been turned off in the UniPrint Client settings). The expectation is within 2000ms. Little will affect this time period with the exception of extremely high CPU and Disk I/O contention. If this occurs, then it's likely that server sizing estimations are incorrect for the number of users and types of applications being used. In particular, the server's CPU may be underpowered or overworked.

UniPrint Client

The UniPrint Client checks for jobs that have been transferred to the client workstation every 1000ms. Testing the speed at which the UniPrint client processes print jobs can be done without making a Terminal Server connection.

² The same application, document, and driver optimizations should be used for each test.

1. On the client workstation start the UniPrint Client from the start menu.
2. Right-click on the **UniPrint** icon in the workstation's icon tray and select **Settings** from the context menu.
3. Click on the **General** tab and configure UniPrint to **Preview with UniPrint**
4. Click on the **Advanced** tab and make note of the **Local Spool Path**. This is the location where the UniPrint Client expects to find UniPrint print jobs. Click **OK** to save the changes.
5. Rename any existing PDF file with a .PPF extension, e.g., change testfile.pdf to testfile.ppf (preferably, use one of the UniPrint jobs that you may have trapped or saved in earlier experiments).
6. Copy the renamed file to the directory identified as the Local Spool Path above.

The UniPrint Client should begin processing the file within 1000ms and then display it in the UniPrint preview window. Measure the time it takes the client to display the PDF file in the preview window. This time period should be equal to or faster than manually opening the file in Acrobat Reader. Alternatively, the client can be configured to use Acrobat Reader for previews. It may be useful to measure the difference in load times between UniPrint's native previewing engine and Acrobat Reader. The UniPrint preview engine is much smaller and should have faster load times.

UniPrint Gateway

The UniPrint process responsible for initiating the transfer of Gateway jobs uses two criteria, in conjunction, which are fully configurable through the UniPrint Configurator:

1. Spool Server Buffer Flush Threshold (default is 30K)
2. Connection Interval to Spool Server (default is 10 seconds)

The Spool Server Buffer Flush Threshold value specifies the maximum amount of cumulative print job data before initiating a transfer. The default threshold of 30K is almost always reached by the first few print jobs and is therefore the principal initiator of the job transfer – this is an event driven process. The Connection Interval to Spool Server setting ensures that print job transfers are initiated every ten seconds regardless of the buffer size. This ensures that any job data not exceeding the buffer threshold are transferred. If only extremely small jobs are being created at any time (for example, a dozen 2K jobs), then it is reasonable to lower the buffer threshold to speed up the initiation of the job transfer.

UniPrint Print Server

Each local and distributed UniPrint Print Server polls the Spool Server for print jobs using FTP. A print server collects the PDF print jobs for which it is responsible and spools them using the native driver of the destination print device. The default polling interval is five seconds and fully configurable using the Print Server Configurator. The maximum Connection Interval to Spool Server setting is 60 seconds.

To determine if the UniPrint Print Server is collecting its print jobs according to the configured time, follow the steps below.

1. Open Services in Control Panel and stop the **UniPrint Print Server – File Management service**. This will also stop the dependent services.
2. Start a session to the Terminal Server and print to a Gateway printer. Be sure to select a printer managed by the print server on which you stopped the services.
3. Your print job will be spooled on the Terminal Server then transferred to the Spool Server. Since the Print Server services have been stopped, the print job will remain in: %systemdrive%\Inetpub\ftproot\UniPrint\PrintJobs\\
4. From the spool server computer, open the above folder and verify that the print job is sitting in the above directory on the Spool Server (You will have to open the folder using a UNC path, e.g., \\Spool_Server\C\$\Inetpub \ftproot\UniPrint\PrintJobs\)
5. To measure the Polling Cycle Time, observe the print job in the spool server folder above and start the **UniPrint Print Server – Print Job Maintenance**. Within the configured interval time, you will see the .PCF file renamed with a **_out** extension. This renaming process indicates that the print server has successfully found the print job and initiated the file transfer. The length of time between starting the Print Job Maintenance service and the renaming of the cached .PCF file should match the configured interval time.

Keep in mind that Print Server polling is done over the network using FTP. The FTP request is to initiate the file transfer only consumes a small amount of bandwidth. However, the polling rate will be affected by bandwidth contention and could vary by as much as 50 - 75% (+2.5 to +3.75 seconds)

Job Transfer Time

Job Transfer Time is how long a UniPrint print job takes to move from one UniPrint component to the next over the LAN, WAN, or Internet. Strictly speaking, job transfer time is directly related to the size of the print job and the performance of the network over which the job is being transferred. Network throughput is the amount of data transferred from one place on the network to another or the amount of data processed in a specified amount of time. UniPrint's job transfer performance can only be measured in comparison to other printing methods such as auto-created printing or traditional network printing.

UniPrint differs from these printing methods in two significant ways:

1. PDF file format and
2. Transfer protocol.

Network throughput can vary greatly if file format or transfer protocol vary.

PDF files are typically smaller than PCL, Postscript, and EMF files – this has a direct impact on network throughput. In the case of UniPrint Server, file format may be the only significant difference affecting job transfer performance compared to other methods using the same transfer protocols. UniPrint Gateway uses FTP to move print jobs compared with traditional network printing which uses Microsoft's Server

Message Block (SMB) protocol. FTP is faster (achieves better throughput) than SMB such as file size and network conditions. SMB sends approximately 35 times more packets than FTP during its initial negotiation process and its performance varies significantly between different Windows operating systems. Overall, FTP outperforms SMB. In high volume printing environments, the increased overhead of the SMB protocol can have a significant impact on the network, and consequently printing performance.

At a different level, but no less significant, is the architecture of the printing solution. For example, auto-created printing can route jobs to network printers in two ways: directly to a print server, or via the client workstation to the print server, both using SMB. The former sends the same job across the network twice, the latter sends it across the network three times. Throughout a print job's travels, its file format changes as it passes from one component to the other and consequently so does its impact on the network. The tables below demonstrate the print job flow from Terminal Server to physical printer, the file format changes that occur along each transfer point and the transfer protocol used. When the file format changes so does the file size.

Auto-created printing job flow (print server mapped and client mapped) →

| | | | | | | | |
|---------------------------------------|----------------------------------|--------------------|----------------------------|--------------|----------------------|------------------|---|
| Terminal Server (Print Server mapped) | PCL5/EMF (1000K over SMB) | | | Print Server | RAW (6000K) over SMB | Physical Printer | Total data transferred (7000K) Time = 8 sec. |
| Terminal Server (Client mapped) | PCL5/EMF (1000K) over ICA or RDP | Client Workstation | PCL5/EMF (1000K+) over SMB | Print Server | RAW (6000K) over SMB | Physical Printer | Total data transferred (8000K) Time = 11 sec. |

UniPrint job flow (UniPrint Server and UniPrint Gateway) →

| | | | | | | | |
|------------------------------------|---------------------------|--------------------|---------------------|--------------|---------------------|------------------|--|
| Terminal Server (UniPrint Server) | PDF (66K) over ICA or RDP | Client Workstation | RAW (858K) over SMB | Print Server | RAW (858K) over SMB | Physical Printer | Total data transferred (1782K) Time = 5 sec. |
| Terminal Server (UniPrint Gateway) | PDF (66K) over ICA or RDP | Spool Server | PDF (66K) over FTP | Print Server | RAW (858K) over SMB | Physical Printer | Total data transferred (990K) Time = 2 sec. |

Comparing the un-shaded rows above, there is a clear difference in the amount of data being transferred using auto-created client printers versus UniPrint client printers (UniPrint Server). The clear difference is the size of the job being transferred over the ICA stream – In general, the equivalent PDF file is 15 times smaller. The transfer time of the EMF file is understandably higher given the larger file size. However, if file size were the same, then network throughput should also be the same given the same network conditions and the fact that the same transfer protocols are being used to move the print job from Terminal Server to printer.

Comparing the shaded rows above, there is also a significant difference in the amount of data being transferred, and greater throughput is achieved with UniPrint

Gateway despite the extra network hop. The significant differences with this comparison are file size and transfer protocol – PDF and FTP versus EMF and SMB.

UniPrint Server The above job transfer comparison helps to provide printing expectations of UniPrint versus standard Windows printing (EMF-based printing). Both use the same transfer protocol, ICA or RDP. Therefore the only difference between the two solutions is file size.

UniPrint Gateway Similar to UniPrint Server, if the size of the file being transferred is the same (that is an EMF file and PDF file of equal size) then the difference in transfer speed will be related to the transfer protocol, SMB versus FTP. To compare transfer times with UniPrint Gateway, save a UniPrint PDF print job on the Terminal Server and transfer it to the UniPrint Spool Server using the provided Connection Validator or using command-line FTP. Measure the time it takes to transfer the file. Repeat the experiment using a Windows mapped drive from the Terminal Server to the same UniPrint Spool Server. A Windows mapped drive uses SMB, the same protocol used for transferring EMF print jobs. Measure the time it takes to transfer the same file using SMB. The time differences should reflect the overhead of SMB compared to FTP.