

Print Tracking and Control for Windows NT/2000/XP/2003

White Paper

This document discusses the subject of print tracking and control, specifically as it applies to Windows NT/2000/XP/2003

The purpose of the document is to provide some basic guidelines and requirements for a print tracking and control system

Also included is a high level overview of the Windows print architecture and the terminology associated with it.

PRINT TRACKING AND CONTROL

As with other quota management software, a print tracking and control system (also called quota management) is used to control and limit a valuable resource, namely wear and tear on printers, paper and toner. Just like disk space, printers are abused, and many thousands of reams of paper are wasted annually, in average corporations. This is a large cost, and is hard to control without the proper data and tools.

Other networks such as Novell NetWare and Unix have software to control and track printing built in. Windows however, does not, and leaves an MIS manager quite helpless when it comes to controlling and tracking printing.

The features provided by Windows for print tracking and control are very rudimentary and include only auditing type options, with no control ability whatsoever. The auditing available is centered on print events, such as job added, job deleted and job paused. These events get logged to the event viewer. Just like any other system resource, Windows provides security control for access and management of printers.

One major fault with the printer auditing provided by Windows is that it does not count job pages for any print job which is printed from a non Windows NT/2000/XP/2003 client (such as Windows 95/98, Mac or Unix). This makes it useless for many corporations whose desktop machines are non-Windows. Microsoft does not publicize this fact, and most MIS managers find it out the hard way.

The only option left to an MIS manager who wants to control printing is to purchase a costly printing control system. These solutions are hardware based (involve custom print servers and hardware accessories) and are very costly. This puts them out of reach for most MIS managers who want to cut a few hundred dollars per week in excess printing.

This is where print tracking and control software comes in, providing a real (affordable) solution for tracking and control.

At a high-level, the functions that a print tracking and control system should perform can be broken down into three areas:

Tracking

- Ability to show the number of pages printed by printer, by user.
- Ability to track printing to specific printers (such as more costly color printers)

- Ability to keep history records for print traffic by user and by printer (for determined heavy usage periods).

Control

- Ability to set a quota on a user, specifying either a dollar figure, or page count.
- Denial of service when this quota is exceeded by a user.
- Ability to set quota's differently for various printers. (I.e. allowing fewer pages on more costly color printers).
- Denial of service on jobs which exceed a maximum size threshold. This can be used to prevent users from jamming up print queues during heavy usage periods.
- Denial of service on jobs which are printed from specific applications. This can be used to prevent users from printing documents from the Internet, or from specific programs such as graphics programs.
- Ability to control what type of action to take when any restriction has been activated.
- User notification (through pop-up messages or printed pages) when their print job has been denied.
- Ability to either Delete, or leave a job when a restriction has been activated.
- The ability to perform all restrictions on specific users on specific printers.
- Ability to pause all jobs which enter a specific printer queue from a specific user(s).
- The ability to perform all restrictions on specific printers.

Reporting

- Ability to perform reports on printing by user, printer or user group. These reports are typically used to charge back departments for printing costs.
- Ability to spot gross offenders by sorting a list of users by pages printed.
- Should allow generic reporting tools to be run against the database backend, as MIS managers have specific tools that they are already familiar with, without having to learn new ones.

TECHNICAL REQUIREMENTS

This section lists specific requirements, benchmarks and test cases that can be used to evaluate and measure any print tracking and control system for Windows NT/2000/XP/2003

- The print management system should seamlessly integrate into the Windows Spooler architecture, ideally requiring no replacement of existing spooler components (e.g. Print Processors, Print Monitors, Print Drivers). This requirement exists to ensure compatibility with all types of processors, drivers, monitors and other components of the Spooler.
- A print management system should be compatible with all connection types, specifically Parallel, Serial, LPR, JetDirect (and all other Network protocols).

The test case to ensure this is to set-up at least one printer using each of these connection methods, and ensure the proper operation of the print device under heavy load situations, such as those described above.

- The print management system should not interfere with paper-out notifications from the printer, and other more advanced notifications (such as those provided by bi-directional printers, i.e. low toner and manual paper feed requests).

The test case to ensure proper operation is to create a condition where a printer runs out of paper, or has a paper jam, and ensure that the proper pop-up messages get sent back

to the client. The proper pop-up messages can be examined by first running the test case without the print management software installed.

- A print management system should impose as little overhead as possible on a print server. In a large system, Windows servers are designed to handle as many as 64 printers, all on one print server. This is a large amount of throughput, when all of these printers are being heavily used. The print management software should not slow-down or stall documents during heavy load situations.

A test case can be set-up to benchmark the overhead imposed by such a system. The test case involves printing a large document (1-5Mb) to a print server with and without the print management software loaded and then timing the total time to pass through to the printer. The time can be measured from the job being sent from the client, until the job is completely printed. The overhead imposed by the print management software should be less than a 10% delta.

- Each print queue should act completely independently of all other print queues. The Windows Spooler is multi-threaded, and so the print management system must also be multi-threaded to prevent cross queue job blocking.

The test case to ensure this requirement involves multiple documents being printed from multiple clients to multiple print queues (all at the same time). Specifically, (at least) three clients print documents of considerable size (1-5Mb) to three different queues (all at the same time). During this test, the large document should be printed multiple times by each of the three clients to ensure that there are documents spooling while others are being downloaded to the printers. This test case simulates real-time heavy load on a print server. The purpose of this test case is to ensure that such a system does not slow down the multi-tasking performance of the Windows spooler

- No “ready to print” document should be blocked by another document which is still spooling.

The test case to ensure this requirement involves multiple large documents (1-5Mb) being spooled from one client, while another client spools small documents (1Kb). In this test, the smaller documents should not get blocked by the larger documents. This test will ensure the proper handling of spooling documents, and ensure that the print management system is not blocking printing while it is tied up on a document that is spooling.

- The print management system should properly identify the page counts as output on the actual print device itself.
- Print jobs should always appear in the same form as they would without the print management software installed. In other words, the print management software should never alter the contents of a print job.
- The print management system should properly identify jobs from all types of Windows clients including Windows NT/2000/ME/XP workstation, Windows 9x, Mac, Unix.
- Full compliance to Windows security standards are required. Specifically, the Access Control List for a given job should never be modified in any way. Additionally, the user list should be properly read from the current PDC if the server is participating in a domain.
- Full support for PDC Users and User groups should be present.

OVERVIEW OF THE WINDOWS PRINT SPOOLER ARCHITECTURE

Print Spooler

The primary component of the printing system is the print spooler. The print spooler is a Windows executable that manages the printing process. Management involved loading the correct print drivers, conversion of GDI function calls to some type of printer language, allocating storage in a disk file for the document, and management of multiple jobs being sent to a printer at one time.

The spooler is loaded at startup and runs until the system is shut down. Windows provides a rudimentary print manager to configure the spooler, however when this program is shutdown, the spooler continues to run.

In a network, the Spooler will accept print jobs from other spoolers. When a printer has been shared, the client spooler will communicate with the server spooler, and transfer the document to be printed. It should be noted that the transfer takes place after the print driver on the local machine has rendered the print data into raw device commands.

Print Processor

Once the spooler determines that a job should be printed, it calls the print processor. The print processor is a Windows DLL that reads and converts the generic graphics output from a Windows application into DDI calls.

Graphics Engine

A Windows DLL whose role is to convert print processor output into calls to device driver functions. The device driver processes these calls and converts them into printer specific device commands that the printer can process. These commands are usually Postscript, PCL or raw text.

Monitor

Once a device driver has translated a journal file into printer specific device commands (Postscript, PCL, etc), the data is passed back to the spooler. The spooler then sends these printer specific commands to a monitor (also called a port monitor). A monitor is a Windows DLL that passes the printer specific device commands over the network, through a parallel port, or through a serial port to the device.