

JULY 1991

Extended Print Services: Obtaining Information for Queue Management

Rich Lee
Senior Consultant
Systems Engineering Division

The print services series of AppNotes examines the technical aspects of NetWare printing, providing configuration information, performance test results, and other tools for handling network print service problems. The April AppNote discussed the effect of printing ports, configurations, and platforms on print service performance, while the June AppNote introduced methods for measuring performance in the high-end printing environment. In this AppNote, we present some tools for collecting information that will help you manage busy print queues.

Copyright © 1991 by Novell, Inc., Provo, Utah. All rights reserved.

As a means of promoting NetWare AppNotes, Novell grants you without charge the right to reproduce, distribute, and use copies of the AppNotes, provided you do not receive any payment, commercial benefit, or other consideration for the reproduction or distribution, or change any copyright notices appearing on or in the document.

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of these Application Notes (AppNotes) or of any of the third-party products discussed in the AppNotes. Novell reserves the right to revise these AppNotes and to make changes in their content at any time, without obligation to notify any person or entity of such revisions or changes. These AppNotes do not constitute an endorsement of the third-party product or products that were tested. Configuration(s) tested or described may or may not be the only available solution. Any test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state, or local requirements. Novell does not warranty products except as stated in applicable Novell product warranties or license agreements.

Contents

Introduction to Queue Management	17
Review of Previous Printing Series AppNotes	17
Test Variables Revisited	18
New Derived Measurements for Queue Management	19
Queue Adding Time	19
Queue Active Time	20
Finish Printing Time	20
Tools for Queue Management	20
Average Printing Time	20
Average Queue Adding Time	21
Average Queue Active Time	22
Average Finish Printing Time	22
Average Printing Time Restated	23
Gathering Queue Usage Information	23
The QSCAN.EXE Program	24
Comparing Average Print Times	24
Design Study	24
Short Cut Calculation	25
Average Adding Rate	25
Average Document Size	26
Average Queue Adding Time	26
Finishing the Calculation	26
Applying the Formulas	27
New Test Results	27

Introduction to Queue Management

A typical LAN installation starts out small—a couple of users and a printer or two. As the LAN grows, adding more applications and more users, the difficulty in managing the network grows exponentially. Performance suddenly becomes an issue. Configuration maintenance is an ongoing difficulty. As applications proliferate, network supervisors find that seemingly beneficial software upgrades pose hidden threats as well. And before you know it, the network printers become deluged with users' print jobs.

But print queue overload usually doesn't happen overnight. Like a complex puzzle, printing dilemmas must be solved one piece at a time. This AppNote discusses tools for print queue management: the process of measuring activity in the queue and obtaining a baseline measurement of significant queue events. It also gives possible directions for using these tools for future queue management. A review of some of the highlights from previous Appnotes will help put our current discussion into the proper perspective.

Review of Previous Printing Series AppNotes

In the "NetWare Print Service Performance" AppNote published in April 1991, we reviewed various Novell print server configurations (PSERVER.VAP, PSERVER.NLM, and PSERVER.EXE) and presented results from a battery of performance tests on each configuration. From these tests, all based on a low-end scenario of only one job in the print queue, we hoped to see what factors affected printing performance the most. We measured performance using two different yardsticks: (1) printing rate, or the number of bytes per second (BPS) processed by the printer; and (2) printing performance, or the number of pages per minute (PPM) put out by the printer.

We found that, more than any network-related factor, the printer port had the greatest effect on the speed of network printing. When connected to a parallel port, both Hewlett-Packard PCL-5 and PostScript printers produced significantly faster results than when connected to a serial port. The second most influential factor was the application used to send the document to the printer. Surprisingly, we observed extremely low performance (PPM) results in all configurations tested.

For the "Measuring Performance in the Extended Print Service Environment" AppNote in June 1991, we set out to measure performance in more typical printing environments with more than one job in the queue at a time. Due to the increased number of variables this entailed, we scaled down our test configurations. Beginning with that AppNote, our basic test configuration consists of only Hewlett-Packard

HPIII printers connected to parallel ports or via the NIC interface. Our workstations and file servers are all 386/20Mhz machines (although eventually we hope to move up to a 486/33MHz server).

With this new test suite, we updated some of our earlier performance results. Several of the results were different—not because print times suddenly changed across the board, but because our equipment and our ability to measure them had improved. The biggest surprise was the fast throughput transmission rate of NIC-based printing with the HPIIIsi.

Throughout the print services series of AppNotes, we will present any updated test results that might be significant.

Test Variables Revisited

As we moved from the low-end of printing to the high-end, we saw the need to dissect and categorize the various events involved in transmitting a document from a workstation to the print queue and ultimately to the printer. In the June AppNote, we defined a number of variables that contribute to overall printing performance, and provided mathematical equations for several others.

Specifically, we defined dependent variables t_0 through t_7 that correspond to the measurements taken in our study of print services. These variables are restated in Figure 1.

Figure 1: The dependent variables we used in measuring print service events.

Variable	Definition
t_0	Start of measurement (document 1 sent to print queue)
t_1	PCONSOLE shows document 1 as "Adding" to the queue
t_2	PCONSOLE shows document 1 as a "Ready" print job in the queue
t_3	PCONSOLE shows document 1 as the "Active" print job in the queue
t_4	First page of document 1 starts to eject at the HP printer
t_5	PCONSOLE clears document 1 from the queue
t_6	Last page of document 1 ejected at the printer (job is finished)
t_7	PCONSOLE shows document 2 as the "Active" print job in the queue

From these "t" variables, we derived some time parameters ("T" formulas) to represent what we felt were significant measurements for our tests. These measurable print service events should prove extremely useful in our ongoing study of allocating print service and printer resources, so we gave them descriptive names as well (see Figure 2).

Figure 2: The derived values we used in measuring print service events.

"T"	Name	Derivation	Description
-----	------	------------	-------------

Value			
T ₀	Start Time	n/a	Beginning of test measurement
T ₁	Dwell Time	t ₁ - t ₀	Interval between Start Time and the time the document starts being added to the queue
T ₂	Adding Time	t ₂ - t ₀	Time required for application to send formatted data to the queue
T ₃	Activation Time	t ₃ - t ₀	Time from start until "Active" status appears in PCONSOLE
T ₄	Throughput Time	t ₄ - t ₀	Time from start until first page starts to eject at printer
T ₅	Clear Time	t ₅ - t ₀	Time from start until the job is cleared from the queue
T ₆	Print Time	t ₆ - t ₀	Time from start until the final page is ejected from the printer

From time to time, we'll need to redefine these variables and time measurements for a specific experiment. For queue management purposes, we must consider application formatting time. Thus, we'll redefine Start Time (T₀) as the point at which the user issues the "print" command from within an application. To simplify the average queue utilization calculation, we'll ignore the t₇ variable and concentrate on measurements for each document in the queue separately.

New Derived Measurements for Queue Management

As a foundation for our discussion on gathering queue management information, we'll divide the entire network printing process into three more manageable components: Queue Adding Time, Queue Active Time, and Finish Printing Time. We can derive these components from the "t" variables defined above.

Queue Adding Time. This measurement, defined as t₂ - t₀, represents the time an application needs to format a document for the printer and fully add the job to a network print queue. Technically, this is the same as the Adding Time described in Figure 2. However, we include "Queue" in the name to distinguish it in the context of queue management.

Queue Active Time. Once the application has fully added a document to a queue, PCONSOLE flags the job "Ready" for the print server. The print server then takes over and sends the job to the printer. During this process, PCONSOLE shows the job as "Active" in the queue. Thus we define Queue Active Time as $t_5 - t_3$, the amount of time the job remains active in the queue. At some point (t_4) during this interval, the first page of the job starts coming out of the printer.

Finish Printing Time. As soon as the print server has finished feeding the job to the printer, PCONSOLE clears the job from the queue. However, pages continue to print at the printer. Finish Printing Time, $t_6 - t_5$, represents this lag time between the removal of the job from the queue and the actual completion of the print job at the printer.

When you tally up the Queue Adding Time, the Queue Active Time, and the Finish Printing Time, you have the total time it takes for a document to be printed from within an application to a network printer. Figure 3 summarizes these three components of total printing time.

Figure 3: New derived measurements for queue management.

Name	Derivation	Description
Queue Adding Time	$t_2 - t_0$	Time required for application to format and add job to the queue
Queue Active Time	$t_5 - t_3$	Time during which the job remains "Active" in the queue
Finish Printing Time	$t_6 - t_5$	Interval between clearing of job from the queue and completion of job at the printer

Tools for Queue Management

Using these derivative times, we can begin to put together a toolkit that LAN supervisors or print service managers can use to determine queue usage on their LAN.

Average Printing Time

In the June AppNote, we introduced a formula for calculating the Average Printing Time for a particular print queue. The result of this calculation indicates approximately how long the *average* person to send a job to a busy print queue can expect to wait for the print job to be completed. We repeat that formula (previously numbered 2-14) as a tentative first tool for print queue management. In doing so, we recognize that hypothetical "average" people certainly get more than their fair share of study, much of which cannot be applied directly to real people in real situations.

In this equation, $T_{\text{aafq}\#}$ (# replaced with 1) represents the average measured time for an application to fully print a document through a particular queue, from start to finish. $T_{\text{aaq}\#}$ is the average time a job stays active in the queue, and N is the number of people using the print queue. Admittedly, $T_{\text{aafq}\#}$ and $T_{\text{aaq}\#}$ are pretty messy-looking variables. Using our newly defined measurements, we can simplify things at least visually, if not conceptually.

As we've already stated, the total printing time for one document is the sum of the Queue Adding Time, the Queue Active Time, and the Finish Printing Time. Thus we can restate *average* printing time as

$$\begin{array}{ccccccc} \text{Average} & & \text{Average} & & \text{Average} & & \text{Average} \\ \text{Printing} & & \text{Queue} & + & \text{Queue} & + & \text{Finish} \\ \text{Time} & = & \text{Adding} & & \text{Active} & & \text{Printing} \\ & & \text{Time} & & \text{Time} & & \text{Time} \end{array} \quad (3-2)$$

From previous studies, we know that printing time depends on a number of factors:

- the types of documents being printed (text only, text and graphics, graphics only, or string)
- the transmittable size of the documents (in bytes)
- the production rate for the printer (in bytes per second)
- the rate of data transmission across the network

For the sake of simplicity, we have ignored the potential "Ready" time between the time a print job finishes adding to the queue and the time the print server polls the queue for the next job. This can be as long as 15 seconds, depending on when the job finishes adding.

With this in mind, let's look at how to calculate the averages in equation 3-2 above.

Average Queue Adding Time

Adding Time is not dependent on anything related to NetWare print services. Rather, it depends entirely on the application (this has held true for every application we have tested so far). In theory, you could get a reasonably accurate Average Queue Adding Time by totalling all the Adding Times for each application used to send jobs to the queue and then dividing by the number of print jobs sent. However, this manual calculation would be practical only in highly simplified situations where a few users print repetitive documents from a single application. Ideally, you could develop a print queue usage profile for every user in the queue group and factor this in to the calculation, but that would be extremely difficult. For most real-world situations, especially pre-implementation design studies, an estimate will have to suffice.

A much easier approach is available if you know the average document adding rate (in BPS) for an application. Then all you need is the average size of the documents being printed, and you can calculate Average Queue Adding Time according to the following formula:

$$\begin{array}{l} \text{Avg.} \\ \text{Queue} = \\ \text{Adding} \\ \text{Time} \end{array} = \frac{\text{Avg. Transmittable Document Size}}{\text{Avg. Adding Rate}} \quad (3-3)$$

We'll show an example of this type of calculation later in this AppNote.

Average Queue Active Time

Average Queue Active Time is governed by the average queue active rate (in BPS) and the transmittable document size, as shown in the following formula:

$$\begin{array}{l} \text{Avg.} \\ \text{Queue} = \\ \text{Active} \\ \text{Time} \end{array} = \frac{\text{Avg. Transmittable Document Size}}{\text{Avg. Active Rate}} \quad (3-4)$$

If the average active rate were nearly constant for different document types, the Average Queue Active Time would be independent of application document type and fall completely within the domain of NetWare's control. Such a relationship could indicate the probability of producing a particular job in a given time frame, and it would also greatly simplify our queue management formulas.

However, proposing such a relationship now might be a bit misleading because you cannot use this information to calculate when a job will be produced at the printer. Even if the rate of data transmission across the network were constant (which it is not), printing performance would still be dependent on factors such as network load, applications, and so forth. But it can be a useful tool for long-term queue management.

Average Finish Printing Time

Generally, you don't need to be too concerned with the Average Finish Printing Time. With the HP1100 printer, the Finish Printing Time does not appear to be NetWare dependent, and we suspect the same holds true with other types of printers as well.

Average Printing Time Restated

By substituting equivalent expressions from formulas 3-2 through 3-4 (and leaving out Average Finish Printing Time), we can restate Average Printing Time formula 3-1 as

Of all our test variables, the only one that appears to be NetWare dependent is Throughput Time, which is primarily influenced by the network transmission protocol (Arcnet, Ethernet, or Token Ring). So why do we include Adding Time and Print Time in the Average Printing Time calculation? We include them to factor in some of the inherent contribution by different types of software and printers, thus lending some realism to the average calculation. Users might change software without consulting the LAN supervisor. Often hardware is changed, reconfigured, or shared without concern for the queue group. Whatever the case, the Average Printing Time is useful for the long-term management of queues on the LAN.

Gathering Queue Usage Information

A chart of average queue-based print times might be useful for low-end planning needs. In low-end environments, it is possible to obtain an actual calculated average. To arrive at that average, however, the LAN supervisor would have to do a lot of manual data gathering. And, of course, you would have to do a sufficient number of tests during periods that represent typical printing activity (not when half of the users are in a meeting).

You must also choose your sample queue group carefully. Four users printing from the same application to one queue, with typically only one job in the queue and a fairly constant workload, would not make a good choice for an ongoing study of print services. Unless, of course, that queue group receives some multitasking software that allows each user to initiate two or possibly three print jobs all at once.

The initial data-gathering labor would be work enough for most of us. However, on a LAN nothing remains static for very long. Users become more proficient with applications and print more often, or you might add new print queues or printers. Unless you can foresee these types of changes and compensate accordingly, your measurement would necessarily have to occur sometime after a change in the queue group. Even if you could automate the final reporting format of the information, you would have to review the average printing time calculation at regular intervals due to the unrefined nature of the measurements and the probability that the network printing setup would change over time. Thus, in the low-end, the value of queue usage information is often outweighed by the labor involved in gathering it.

Such a chart would be far too hard to maintain for any needs at the high-end. For high-end systems, the interval between measurements would need to be close together and the number of observations increased for each queue. In the high-end where there is more than one job in a queue, this could mean an extremely laborious amount of queue management. It would take a lot of measuring to get the

numbers to manage with, and pity the poor individual who sits in front of PCONSOLE to do it!

So developing an information gathering mechanism for the LAN is fraught with inherent difficulties. At the low end, the LAN supervisor might find monitoring PCONSOLE for the information a bit boring. At the high end, the LAN supervisor/print manager could hardly be expected to keep up with the PCONSOLE display.

The QSCAN.EXE Program

As a partial solution, a program to automatically collect the needed information has been provided on NetWare. This program, named QSCAN.EXE, captures data concerning your queues and servers. You can use the information it provides in several ways. (You may find other uses besides those presented here.)

Comparing Average Print Times

One basic way of using average print times would be an old versus new comparison: taking the current average print time and comparing it with the previously measured average print time.

$$\begin{array}{r} \text{Current Average} \\ \text{Print Time} \end{array} - \begin{array}{r} \text{Previous Average} \\ \text{Print Time} \end{array} \quad (3-6)$$

The result would indicate any difference since the last time average printing time was measured.

Even if equation 3-6 above should yield a difference great enough to raise an eyebrow, most LAN supervisors would check their subtraction before giving any amount of credence to this measurement. Two measurements alone are probably not significant. In fact, the result of equation 3-6 really only indicates a change since the previous measurement. For long-term management, a histogram of running averages would be more appropriate.

Design Study

As a characterization of queue activity, Average Print Time is just a beginning step in developing a network print services management model. It tells us a little bit about how users utilize the print queue and network print services. Combining this with information about how applications address print services, how the print services operate, and how printers do their task, it is possible to create a design model for new network implementations as well as for evaluating existing installations.

Developing a queue utilization model for the network is based on the belief that the approximations given in

formulas 3-2 and 3-6 are of some value. The following example might help demonstrate the value of Average Queue Adding Time and Average Queue Active Time.

Short Cut Calculation

If you know the average transmittable size of the documents sent to a queue, it is far easier to work with the average adding rate, as shown in formula 3-3, to obtain the Average Queue Adding Time. For instance, if you know or estimate that highly complex Graphics Only documents add at 5 KBPS (kilobytes per second), and that Text Only documents add at 3 KBPS, you could determine the Queue Adding Time for a 100KB transmittable G/O document and for a 300KB transmittable T/O document as follows:

$$\begin{aligned} 300 \text{ KB} / (3 \text{ KBPS}) &= 100 \text{ sec} \\ 100 \text{ KB} / (5 \text{ KBPS}) &= 20 \text{ sec} \end{aligned}$$

The Average Queue Adding Time for these documents is

$$(100 \text{ sec} + 20 \text{ sec}) / 2 = 60 \text{ sec}$$

Of course, users don't always print the same types of documents. If you have a small sample queue group, you can manually determine the frequency of printing different document types and adjust the calculation accordingly. For a large sample group, you're better off culling usage information with QSCAN.EXE.

Average Adding Rate. Figure 4 below shows some adding rates extracted from our latest test results for three sample documents. (The complete results are given under "New Test Results" at the end of this AppNote.)

Figure 4: Sample document adding rates from the latest tests.

Document Type	Size (in bytes)	Adding Rate (BPS)
Graphics Only	154,729	4,421
Text and Graphics	270,451	4,098
Text Only	241,679	2,843

The average adding rate is the sum of the individual adding rates divided by the number of samples observed:

$$(4,421 + 4,098 + 2,843) / 3 = 3,787 \text{ BPS}$$

This number translates into roughly 3.7 KBPS. Again, you could adjust for the frequency of document types printed by the queue users for a more accurate result.

Average Document Size. Similarly, the average document size is the sum of the individual sizes divided by the number of documents:

$$(154,729 + 270,451 + 241,679) / 3 = 222,286 \text{ bytes}$$

Average Queue Adding Time. Taking the average document size of 222,286 bytes and dividing it by the average adding rate of 3,787 bytes per second yields the following Average Queue Adding Time:

$$222,286 \text{ bytes} / 3,787 \text{ bps} = 59 \text{ seconds}$$

This calculated approximation of 59 seconds compares favorably with the actual Average Queue Adding Time observed in our tests. Using the numbers labeled J_2 , K_2 , and L_2 in Figure 6, for example, gives an observed average of:

$$(35 + 66 + 85) / 3 = 62 \text{ seconds}$$

The calculated time and observed time are within a few seconds of each other. Given the margin of error in polling the queues, the calculated time should be sufficiently accurate for predicting Average Queue Adding Time for a queue. Measure the adding rates and document types/sizes on your system to come up with your own Average Queue Adding Times.

Finishing the Calculation. Once you have the Average Queue Adding Time, you can plug it in to formula 3-2 for Average Printing Time. You still need the Average Queue Active Time and the Average Finish Printing Time, however. The easiest way to obtain the Average Queue Active Time is by taking the average of the d_3 values obtained through the QSCAN.EXE program. In QSCAN, d_3 measures the interval from when the job goes "Active" to the time it is cleared from the queue. The Finish Printing Time is the time measured from the Clear time to when the last page is printed at the printer.

Suppose we calculated an Average Queue Active Time of 22 seconds, and an Average Finish Printing Time of 13 seconds. Our approximation for Average Printing Time would then be

$$\begin{aligned} \text{Adding time} &= 62 \text{ sec} \\ \text{Active time} &= 22 \text{ sec} \\ \text{Finish time} &= 13 \text{ sec} \end{aligned}$$

$$\text{Printing time} = 97 \text{ sec}$$

To compare this approximation with our measured results, we take the average of the numbers labeled J, K, and L in Figure 6:

$$\begin{aligned} \text{Printing Time (G/O)} &= 55 \text{ sec} \\ \text{Printing Time (T/G)} &= 119 \text{ sec} \\ \text{Printing Time (T/O)} &= 156 \text{ sec} \end{aligned}$$

$$\text{Avg. Printing Time} = 110 \text{ sec}$$

The difference between the approximation and the measured results amounts to 13 seconds in 110, but then the sample is very small. With a larger sample you would want to explain the margin of error, but for our purposes this method simplifies what would otherwise take a lot of work to calculate. Using QSCAN removes the human factor in sampling, which should reduce the margin of error and provide an increased sample size.

Applying the Formulas

Armed with these formulas and approximations to help you determine the nature of your print queues, you can now estimate potential queue utilization by calculating the Average Printing Time. High Average Printing Times coupled with user requests for greater network printing speed are a sure indicator that you need to re-evaluate your print service configurations, applications, and peripherals.

New Test Results

The tables below shows our latest test results using WordPerfect 5.1 to print three test document types. Figure 5 shows the numbers obtained when we used an HP LaserJet IID printer connected via a parallel port. Figure 6 shows the results using the HP IIIsi printer connected directly to the network via the NIC interface. The column numbers and letters A, B, C, J, K, and L are for ease of referring to these numbers in the future.

Figure 5: WordPerfect 5.1 test results with an HP LaserJet IID on a parallel port.

Document Type (1)	Document Size (2)	Avg. Total Printing Time (3)	Avg. Queue Adding Time (4)	Avg. Queue Active Time (5)	Avg. Finish Printing Time (6)
Graphics Only	154,729	A 91 sec	A ₂ 31 sec	A ₃ 45 sec	A ₄ 13 sec
Text & Graphics	270,451	B 245 sec	B ₂ 52 sec	B ₃ 42 sec	B ₄ 8 sec
Text Only	241,679	C 297 sec	C ₂ 43 sec	C ₃ 36 sec	C ₄ 18 sec

Figure 6: WordPerfect 5.1 test results with an HP IIIsi printer connected directly to the network cable via a NIC.

Document Type (1)	Document Size (2)	Avg. Total Printing Time (3)	Avg. Queue Adding Time (4)	Avg. Queue Active Time (5)	Avg. Finish Printing Time (6)
Graphics Only	154,729	J 55 sec	J ₂ 35 sec	J ₃ 5 sec	J ₄ 13 sec
Text & Graphics	270,451	K 119 sec	K ₂ 66 sec	K ₃ 25 sec	K ₄ 15 sec
Text Only	241,679	L 156 sec	L ₂ 85 sec	L ₃ 18 sec	L ₄ 10 sec