



Planning and Deployment Guide

Version 2.5.1



Copyright

2016. Dragon Medical Network Edition, 2.5.1.

This material may not include some last-minute technical changes and/or revisions to the software. Changes are periodically made to the information provided here. Future versions of this material will incorporate these changes.

Nuance Communications, Inc. has patents or pending patent applications covering the subject matter contained in this document. The furnishing of this document does not give you any license to such patents.

No part of this manual or software may be reproduced in any form or by any means, including, without limitation, electronic or mechanical, such as photocopying or recording, or by any information storage and retrieval systems, without the express written consent of Nuance Communications, Inc. Specifications are subject to change without notice.

Copyright © 2002-2016 Nuance Communications, Inc. All rights reserved.

Nuance, ScanSoft, the Nuance logo, the Dragon logo, Dragon, DragonBar, NaturallySpeaking, NaturallyMobile, RealSpeak, Nothing But Speech (NBS), Natural Language Technology, Select-and-Say, MouseGrid, Vocabulary Editor, are registered trademarks or trademarks of Nuance Communications, Inc. in the United States or other countries. All other names and trademarks referenced herein are trademarks of Nuance Communications or their respective owners. Designations used by third-party manufacturers and sellers to distinguish their products may be claimed as trademarks by those third-parties.

Disclaimer

Nuance makes no warranty, express or implied, with respect to the quality, reliability, currentness, accuracy, or freedom from error of this document or the product or products referred to herein and specifically disclaims any implied warranties, including, without limitation, any implied warranty of merchantability, fitness for any particular purpose, or non-infringement.

Nuance disclaims all liability for any direct, indirect, incidental, consequential, special, or exemplary damages resulting from the use of the information in this document. Mention of any product not manufactured by Nuance does not constitute an endorsement by Nuance of that product.

Notice

Nuance Communications, Inc. is strongly committed to creating high quality voice and data management products that, when used in conjunction with your own company's security policies and practices, deliver an efficient and secure means of managing confidential information.

Nuance believes that data security is best maintained by limiting access to various types of information to authorized users only. Although no software product can completely guarantee against security failure, Dragon Medical Network Edition software contains configurable password features that, when used properly, provide a high degree of protection.

L-3870

We strongly urge current owners of Nuance products that include optional system password features to verify that these features are enabled! You can call our support line if you need assistance in setting up passwords correctly or in verifying your existing security settings.

Published by Nuance Communications, Inc., Burlington, Massachusetts, USA

Visit Nuance Communications, Inc. on the Web at www.nuance.com.

Visit Nuance Healthcare on the Web at www.nuance.com/healthcare.

Contents

Chapter 1: Preparing to install Dragon Medical Network Edition	1
Introducing Dragon Medical Network Edition	2
Audience for this guide	2
Introduction	2
Who works on a Dragon Medical Network Edition network	2
Planning an installation of Dragon Medical Network Edition	3
Chapter 2: Network requirements for Dragon Medical Network Edition	5
Dragon Medical Network Edition Overview	6
Dragon Medical Network Edition components	7
Nuance Management Server hosted on-premise	7
Nuance Management Server components	8
Profile Optimizer components	8
Dragon Clients	8
Master user profiles directory	9
Nuance hosted servers	9
Determining space requirements for user profiles	10
Understanding master user profiles and local cache user profiles	10
Determining storage space required for master user profiles	12
Determining storage space required for local cache user profiles on client machines	15
Deciding how often to update acoustic and language models	17
Understanding ACO and LMO processes	17
Defining ACO/LMO processes and their impact on recognition	17
Factors in how often to run ACO processes	18
Factors in how often to run LMO processes	20
Factors that have no effect on regularly scheduled processes	20
Determining number of Speech Nodes required for optimizations	20
Determining number of Speech Nodes required to process acoustic and language models	21
Recommendations about number of users	24
Estimating storage requirements for the Nuance Management database database	25
About Nuance Management Server (NMS) and NMS SQL database	25

About disk space for server installations and logs	26
Recommended network and switch settings	27
Network Interface Card (NIC) settings	27
Network speed	27
Using Network traffic switch for load balancing	27
Storage hardware requirements for master user profiles	28
Ports to open for clients, servers, and hardware firewalls	29
Installing the product without upgrading from a previous version of the product	29
Upgrading from a previous version of the product	29
Ports to open when using an on-premise NMS	30
Chapter 3: Configuring DM Network Edition networks that use physical servers	33
Using physical servers in a DM Network Edition network	34
Configuring a DM Network Edition network for up to 100 users	34
Configuring a DM Network Edition network for 100 to 1000 users	34
Configuring a DM Network Edition network for more than 1000 users	35
System requirements for small networks with physical servers	39
System requirements for medium networks with physical servers	40
System requirements for large networks with physical servers	42
System requirements for master user profile computers	44
Hardware and software requirements	44
Optional hardware	45
Optional software	45
Summary: System requirements for DM Network Edition networks that use physical servers	46
System requirements for databases, the NMS Server, and the NMC Center	46
System requirements for Speech nodes, Master User Profiles Host, and the Dragon client	47
Summary: Configuration options for DM Network Edition networks that use physical servers	49
Chapter 4: Configuring DM Network Edition networks that use virtual servers	51
Using virtual servers in a DM Network Edition network	52
Using Virtual Machines for the NMS Server, and databases	52
Using Virtual Machines for Profile Optimizer Speech Nodes	52
Tested virtual machine desktop configurations for Dragon and VMWare	52
Special Requirements when running Dragon in a virtual environment	53
System requirements for small networks with virtual servers	54
System requirements for medium networks with virtual servers	55

System requirements for large networks with virtual servers	58
Chapter 5: System requirements for the Dragon client	61
System requirements for the Dragon Client	62
Hardware Recommendations	62
Software Requirements	63
Supported Virtual Desktop Infrastructure Applications and Terminal Emulators	65
Supported non-EHR Applications	65
Chapter 6: Network bandwidth requirements on DM Network	
Edition networks	66
Network Bandwidth Requirements	67
Recommendations for Network Bandwidth	67
Considerations for Low Network Bandwidth Configurations	67
Considerations for Network Bandwidth in Citrix Environments with vSync	68
Estimating Dragon Response Times When user Actions Access NMS Server	69
Approximate Sizes of Data Transferred Over DM Network Edition Network and Effect on Network Load	70
Index	76

Chapter 1: Preparing to install Dragon Medical Network Edition

Introducing Dragon Medical Network Edition	2
Planning an installation of Dragon Medical Network Edition	3

Introducing Dragon Medical Network Edition

Audience for this guide

This guide is designed to help system and network administrators plan the network for deployment of the product network. The guide includes details about hardware, software, and network requirements for server-side and client-side components. For information about how to administer the product using the Dragon client, see the product's administrator guide.

Introduction

Dragon Medical Network Edition provides a central server (the *Nuance Management Server*) to manage multiple local or remote speech recognition client machines where users dictate. DM Network Edition takes advantage of distributed processing across servers and workstations for greater efficiency.

The administrator can manage the entire system from the *NMS Server*.

The *Planning and Deployment Guide* takes you through the steps to:

- Plan a deployment of a *Dragon Medical Network Edition* network.
- Determine the number of servers and other equipment you will need.
- Prepare equipment for the DM Network Edition installation.

Nuance also provides a separate *Installation Guide* that takes you through the process of installing the various server and client components, connecting them to a database that stores your organization's data, and installing/setting up *Dragon Clients* that work with both the *NMS Server* and the other servers and components of the network.

Who works on a *Dragon Medical Network Edition* network

- **System Administrators** — Manage the entire system through the *Nuance Management Center*.
- **Users or Healthcare providers** — Dictate using *Dragon Clients*.

Planning an installation of Dragon Medical Network Edition

Before you begin the installation, you should evaluate your own system installation skill set. If you do not have all of the required skills, you should have someone on standby that can help you with this installation:

- Create a network domain/user account with full read/write access rights across all servers
- Create databases with SQL Server
- Set up backup plans for SQL Server
- Create and securely administer a Windows share
- Set Windows user rights and directory permissions
- Set up and configure Internet Information Services (IIS)
- Securely administer IIS (if using web server for master user profiles)
- Order, receive, and install SSL certificates in IIS (if using secure web server for master user profiles)
- Edit XML configuration files
- Manage Windows Services
- Configure a RAID array

You can choose to create the *NMS Server* database and backup dump device yourself during the installation process, but in many configurations, the Dragon Medical Network Edition installers set up the SQL Server databases and dump devices for you. Later, the servers carry out automatic backups of selected data and retain those backups for particular lengths of time; however, because you might need to back up additional data or retain the data longer, database backup planning skills are essential. For more information on backups in DM Network Edition, see the DM Network Edition *Installation Guide* or the DM Network Edition Administrator guide.

Chapter 2: Network requirements for Dragon Medical Network Edition

You can configure your Dragon Medical Network Edition network using the *Nuance Management Server* in several ways.

Dragon Medical Network Edition Overview	6
Dragon Medical Network Edition components	7
Determining space requirements for user profiles	10
Deciding how often to update acoustic and language models	17
Determining number of Speech Nodes required for optimizations	20
Estimating storage requirements for the Nuance Management database database	25
Recommended network and switch settings	27
Using Network traffic switch for load balancing	27
Storage hardware requirements for master user profiles	28
Ports to open for clients, servers, and hardware firewalls	29

Dragon Medical Network Edition Overview

You use the *NMS Server* to manage multiple Dragon clients from a single central server.

The *Nuance Management Center* allows you to change multiple settings on the *NMS Server*, and apply the settings to many Dragon client workstations in a facility.

You can take cross-network action on *Dragon* clients from a central *NMS Server*:

- Create multiple sites where *Dragon Clients* are installed and assign the same backup location, master user profile storage location, and amount of archive storage/playback space to all users on one site in a single stroke.
- Grant, revoke, and manage all dictating user and administrator (**NMS Administrator**) licenses, including viewing the number/type of licenses used/still available.
- Create groups of users and then assign entire groups to one or more sites, all at one time.
- Search through groups of users for those associated with a particular site.
- View/access *Dragon* logs on every *Dragon Client* workstation.
- Initiate, schedule, and monitor optimization processes—monitoring each process for active versus idle state, start time, completion time, duration, and priority.
- Audit session events across the entire network (who logged on when and other related events) for a particular slice of time or particular user.
- Display a complete list of all downloaded updates and approve or not approve installation of them—all from a single window.
- Create user accounts for users and then use simple software tools provided to quickly upgrade old user profiles, then associate upgraded profiles with new user accounts.
- Receive messages for command updates, medication updates, software updates, notifications, and license expiration notifications in a single location.
- Approve command and medication updates for download.
- Choose to encrypt all patient data for all dictating users in your organization.
- Choose to upload speech data to Nuance's Research department for particular users or all users in your organization.

These features help you to readily manage a network of *Dragon Clients*.

Dragon Medical Network Edition components

Dragon Medical Network Edition is made up of software modules that include clients, servers, databases, user profile directories, and web services that integrate with one another.

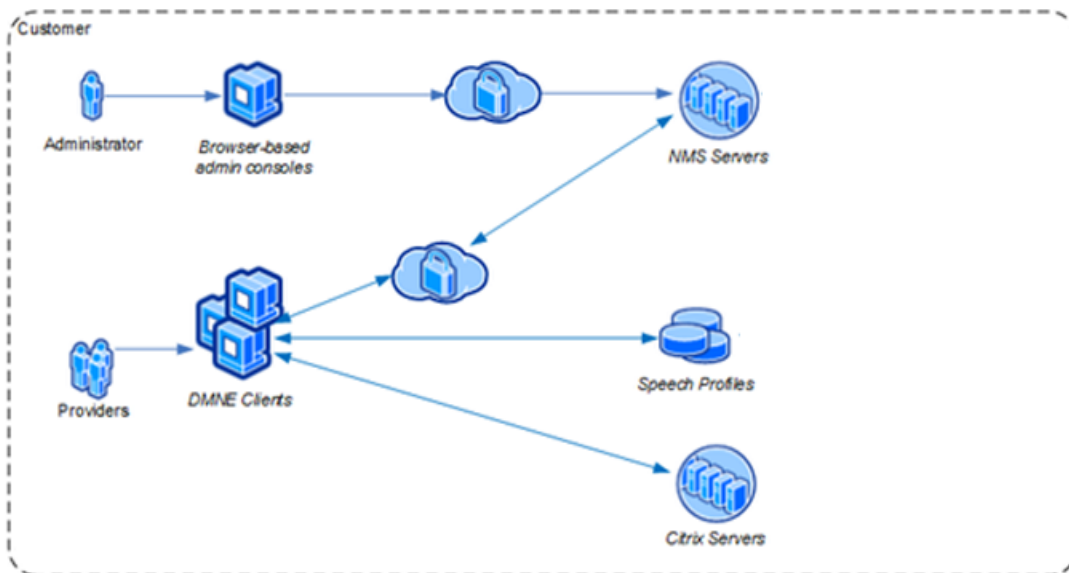
The modules are in these groups:

- *Nuance Management Server* components
- Speech Nodes
- *Dragon Clients*
- Master user profiles directory (on any server of the network or on its own machine)
- Nuance hosted servers

Nuance Management Server hosted on-premise

If your end users do not need to share profile customizations with Nuance cloud-based applications, you install the Nuance Management Server (NMS) components locally. This configuration is referred to as "NMS hosted on-premise."

Nuance Management Server hosted on-premise deployment



- Text and Graphics commands / auto-texts
- Custom words
- Custom Command Sets

For example, if an administrator creates an auto-text in the Nuance Management Center, users can use that auto-text in Dragon Medical Network Edition, Dragon Medical Direct, and applications that use Dragon Medical SpeechAnywhere Services.

Nuance Management Server components

The *Nuance Management Server (NMS Server)* components are:

- Nuance also interacts with your organization through this server, using your unique organization ID to identify your server and downloading licenses you have purchased to it.
The *NMS Server* also stores user identification and related administrative information centrally, keeping user profiles available at any site the user travels ("roams") to within the hospital and its affiliate facilities, whether the machines are located in different offices of the same site or at other sites in distant cities.
- **Nuance Management Center** — An interface to not only the *Nuance Management Server* for managing set up of all *Dragon Clients* on the network.
- **NMS SQL Database** — Houses the SQL database that stores *Nuance Management Server* information about your organization, licenses, sites, users, groups, tasks in process, and messages/downloaded package approvals.

You install the *NMS Server* and the *Nuance Management Center* on the same machine; however, you have the option of installing the database on the same or a separate physical server on your network.

Profile Optimizer components

Profile Optimizer components include:

- **Profile Optimizer Speech Nodes** — Modules that carry out the Acoustic Optimization (ACO), Language Model Optimization (LMO), and other tasks, for each dictating user.

Dragon Clients

Dragon Client — Receives and interprets dictation, producing text corresponding to each users speech. The *Dragon Client* sends the text into a target application, such as an electronic health record (EHR) system, and also preserves text and audio in the speech profiles so that the optimizer can process it.

Master user profiles directory

Master User Profile Directory — Central storage area for master user profiles for dictating users. A single central storage location applies to a particular site in DM Network Edition. A master user profile contains the accumulated speech information for a particular user from all the *Dragon* client workstations where that person performed dictation, enveloping the language the user speaks, the accent (if applicable) used, vocabulary or vocabularies, and microphone and/or dictation recording device. As the user dictates, Dragon learns the particular user's speech patterns and integrates them into that individual user profile. These user profiles are called *roaming user profiles*, because their existence at a central location allows speakers to "roam" from location to location and work with the same profile. (These profiles are unrelated to Microsoft Windows user profiles.)

Nuance hosted servers

Nuance hosts servers and services at its headquarters that your DM Network Edition network interacts with:

- **NuanceLink** — Hosted by Nuance on the Internet. Automatically sends various updates to the *NMS Server*.

Behind the scenes web services on the *NMS Server* help it interact with all of these clients and servers as well as carry out its own processes to manage the entire DM Network Edition network.

Determining space requirements for user profiles

User profiles reside within your network.

To determine how much storage space your network requires to manage user profiles for all of your dictating users, you need to first understand the difference between a master user profile and a local cache user profile.

You then calculate how much storage space you need to work with these two types of user profiles, as explained in the subsections that follow:

- *Understanding master user profiles and local cache user profiles*
- *Determining storage space required for master user profiles*
- *Determining storage space required for local cache user profiles on client machines*

Understanding master user profiles and local cache user profiles

A *user profile* is a set of files containing data that *Dragon* uses to help recognize the individual user's speech.

The user profile information is stored on the network (usually on a separate server or workstation machine), so that the person dictating can "roam" from machine to machine (from one examining room to the next or to a laptop for an off-site consultation) and have consistent recognition accuracy on every workstation. Accessing that information over the network can tend to slow down speech recognition, so to make the process more efficient, the first time the user logs into *Dragon* on a client machine, *Dragon* makes a copy of that user's user profile and stores it locally on that client machine. The user profile stored on the network is called the *master user profile* and the copy stored locally is called the *local cache user profile*.

Each time a user logs out of the *Dragon Client*, information from that client's local cache user profile is synchronized with the master user profile. To minimize network bandwidth consumption, not all data from the local cache profile is retrieved, only the new information from the current dictation session. The next time the user logs into the *Dragon Client*, *Dragon* looks for new information that may have come from dictation sessions on other workstations and copies them into the existing local cache, so again it does not need to copy the entire profile.

For planning purposes, you need to determine how much space your *Dragon* network requires to store *master* user profiles. How much storage space is required depends on how many users you

have and other specific data about those users. Plus, you need to allow space for automatic backups that the NMS Server creates, because it stores those backups on the same device where the master profiles reside.

Similarly, you need to determine how much space each *Dragon Client* machine requires to store *local cache* user profiles for all users that dictate on that machine.

Determining storage space required for master user profiles

Gather the data needed to determine how much space should be allocated for storing master user profiles. This data includes:

- How many total users are dictating with *Dragon* throughout the entire organization?
- How many users have an extra vocabulary? You should add an extra vocabulary for each user who practices in more than one specialty, as well as for any user who also uses *Dragon* to dictate email or create administrative documents other than clinical documents, to allocate adequate space for the extra speech information.
- How many users will have more than one audio input device, such as a headset microphone in some locations and a USB microphone in others? How many extra devices does that add up to?
- Expected growth of your organization—How many new users do you anticipate adding to the population of users dictating with *Dragon* in the next year?

Notes about acoustic models and user profiles

- When you upgrade a profile to DM Network Edition 2.5.1, the Acoustic optimizer processes several .dra files. The processed .dra files are larger than the original .dra files.
- When you upgrade a profile to DM Network Edition 2.5.1, the upgrade process copies the full contents of the backup directory. The size of the folder depends on the size of the current folder that was present when the backup was made.

Calculating storage space required for master user profiles

Number of Users Dictating					
Total Number of users (N)	Users with one extra vocabulary (V)	Number of extra audio input device(s) (A)	Number of users Added Annually (P) (optional)	Acoustic Model Optimization Data Storage (M)	Multiply Total by 2.05 to Allow for Backups
Allow 60 MB for each user.	Allow 20 MB for each extra vocabulary (Add another vocabulary when user also dictates outside system into email or office apps).	Allow 60 MB for each extra audio input device (dictation source).	Add 53 MB for each user, allowing each an extra audio input device and an extra vocabulary	Allow 1 GB of space for each user times the number of audio input devices (dictation sources) the user has.	Multiply the total by 2.05 to allow enough storage space for one backup plus an extra 5% for temporary storage of corrupted profiles.
$((N \times 60 \text{ MB})$	$+ (V \times 20 \text{ MB})$	$+ (A \times 60 \text{ MB})$	$+ (P \times 53 \text{ MB})$	$((N + A + P)(1 \text{ GB}))$	$\times 2.05$
	+	+	+	+	$\times 2.05$
TOTAL:					

Example: Calculating storage space required for master user profiles

Number of Users Dictating					
Total Number of users (N)	users with one extra vocabulary (V)	Number of extra audio input device (s)(A)	Number of Professionals Added Annually (P) (optional)	Acoustic Model Optimization Data Storage (M)	Multiply Total by 2.05 to Allow for Backups
Example:					
100 users	12 with one extra vocabulary and 3 with two extra vocabularies	Each user has a headset and a PowerMic II or PowerMic III. First microphone already included in original space; second considered additional.	Planning to add 5 new users in the coming year.	Multiply number of users by 2 because each has 2 microphones; then multiply that total by 1 GB.	Multiply the total by 2.05 to allow enough storage space for one backup plus an extra 5% for temporary storage of corrupted profiles.
N = 100	V = 18	A = 100	P = 5	M = 205 x 1000 MB	Total x 2.05
((100 x 60)	(18 x 20)	(100 x 60)	(5 x 53)	(205 x 1000))	x 2.05
(6000	+ 360	+ 6000	+ 265	+ 205,000)	x 2.05
TOTAL: 217, 625 MB x 2.05 = 446, 131.25 MB (rounded up) = 446 GB					

As you later see, when you install the *NMS Server*, you indicate the name of the machine and path to the master user profiles. Nuance recommends that you place master user profiles on their own RAID array in *Dragon* configurations that have more than 10 users.

RAID refers to Redundant Array of Independent Disks, a technology that combines multiple small, inexpensive disk drives to form an array whose performance exceeds that of a single large and expensive drive. Your server or workstation sees this array of drives as a single logical storage unit.

Determining storage space required for local cache user profiles on client machines

In addition to calculating how much space you need for master user profiles, you need to ensure adequate space on each *Dragon Client* workstation to make a local copy of the master user profile for the client to use when the user dictates.

Notes about acoustic models

- When you upgrade a profile to Dragon Medical Network Edition 2.5.1, the ACO processes several .dra files. The processed .dra files are larger than the original .dra files.
- When you upgrade a profile to Dragon Medical Network Edition 2.5.1, the upgrade process copies the full contents of the backup directory. The size of the folder depends on the size of the current folder that was present when the backup was made.

Calculating Storage Space Required for local user profiles on clients

Number of Users Dictating				
Number of Users Dictating on Client (C)	Number of Users with one Extra Vocabulary (V)	Number of Extra Audio Input Device (s) (A)	Number of Languages > LMO data (L)	Acoustic Model Optimization Data Storage (M)
Allow 60 MB for each user.	Allow 20 MB for each extra vocabulary.	Allow 60 MB for each extra audio input device(dictation source).	Allow 10 MB for each vocabulary to account for language model optimization data	Allow 500 MB of space for each user times the number of audio input devices (dictation sources) the user has.
$((C \times 60 \text{ MB})$	$+ (V \times 20 \text{ MB})$	$+ (A \times 60 \text{ MB})$	$+ (C + V) (10 \text{ MB})$	$+ (C + A)(500 \text{ MB}))$
TOTAL:				
Example: Allocating space for 10 dictating users:				
10 users	2 with one extra vocabulary and 1 with two extra vocabularies	All users have a headset and a PowerMic II or PowerMic III. The first microphone is already included in original space and the second is	Every user has one vocabulary with a model to be optimized, so add 10 MB for every user.	Multiply the number of users by 2 because each has 2 microphones; then multiply that total by 500 MB.

Number of Users Dictating				
Number of Users Dictating on Client (C)	Number of Users with one Extra Vocabulary (V)	Number of Extra Audio Input Device (s) (A)	Number of Languages > LMO data (L)	Acoustic Model Optimization Data Storage (M)
		considered extra.		
C = 10	V = 4	A = 10	L = 14	M = 20 x 500 MB
((10 x 60)	+ (4 x 20)	+ (10 x 60)	+ (14 x 10)	+ (20 x 500))
(600	+ 80	+ 600	+ 140	+ 10000)
TOTAL: 11420 MB or (rounded up) 11.4 GB				

From the *Nuance Management Center*, you can set the amount of disk space to be reserved for archived corrections for user profiles on each workstation running a *Dragon Client*. When the space is full, *Dragon* automatically moves the archived data to the master profile for use by the Acoustic and Language Model Optimizer.

Deciding how often to update acoustic and language models

The accuracy of dictation depends on how well-tuned the master user profile is for a particular user. *Dragon* tunes each user's master user profile with data from his or her unique acoustic model and language model.

Understanding ACO and LMO processes

The Nuance Management Server (NMS) manages the aspect of speech recognition that involves learning each user's individual pronunciation and speaking patterns. The server controls this learning process (called *adaptation*) by scheduling two processes, Acoustic Model Optimization (ACO) and Language Model Optimization (LMO), which run on the *Profile Optimizer Speech Nodes*. After an optimization process completes, the NMS updates that user's master user profile, integrating the ACO and LMO results.

The accuracy of dictation depends on how well-tuned the master user profile is for a particular user. Research has shown that running both ACO and LMO processes regularly reduces the relative word error rate for the user by 10 to 30% and can produce even greater improvements in recognition accuracy for users who speak with an accent.

When the NMS server attempts to start an immediate ACO or LMO task for a user, and a job for that user and task type is already scheduled (either waiting to run for the first time, or in the postponed state), the NMS server updates the job parameters. If the job is in the postponed state, the NMS server sets the job to run as soon as possible, rather than when the postponed state expires.

When you are determining how often to schedule Acoustic Model and Language Model Optimizations (ACO and LMO processes), you take into account several factors. But how often should the server carry out an ACO or LMO for a particular user?

First, let's understand the purpose of each process.

Defining ACO/LMO processes and their impact on recognition

To help you assess how often an ACO or LMO process should occur, let's first look at what each process does:

- **Acoustic Optimization (ACO)** — This task optimizes the model used to recognize the way the user speaks. The model incorporates characteristics such as the accent, dialect, tone of voice, and other aspects of how the speech sounds. When a speaker is new to *Dragon*, you should run

the optimization more often than after the speaker has worked with the product for a while.

- **Language Model Optimization (LMO)** — This task optimizes the language model used to help recognize the patterns of the user's speech. The model incorporates how the user dictates particular words and sequences of words his or her speech frequently employs.

Both of these processing run in the background and strictly on the *Speech Nodes*, so they do not affect the speed of recognition during dictation on *Dragon Clients*.

You can and should run the ACO and LMO processes at different intervals for the same user.

Factors in how often to run ACO processes

The ACO process, by contrast with LMO, is time intensive—it runs for 1 hour to process 1 hour of dictation. *Speech Nodes* spend more than 80% of their time running ACO processes.

Some of the factors to consider when deciding how often to carry out an ACO process are:

- Whether or not a particular user has an accent
- Whether or not the product frequently has trouble interpreting a particular user's speech
- How much dictation a user or group of users generates

To get the most out of the ACO process, you should schedule the process to take place every time the user has dictated 5 hours of audio. A user does not produce 5 hours of audio by working 5 hours, but may produce that much audio over a period of a month or two, while dictating as needed.

The absolute minimum frequency recommended for ACO is once every 6 months. However, if the user dictates 10 hours of audio in those 6 months, the ACO process will incorporate only the last 5 hours of his or her dictation, so in such a case you would want to schedule the ACO every 3 months instead of every 6 months.

The best practice recommendation is to schedule an ACO every week only while the user is new to *Dragon*. After the user has been dictating for a month or has dictated 5 hours of total audio, the user is no longer considered new to dictation. You can then back off to scheduling the ACO once a month or as often as you find that the user dictates 5 hours of new audio.

If the *Profile Optimizer Speech Node* does not find enough speech data for that user in a given month, the node terminates the process and waits until the next scheduled ACO.

If a professional uses more than one audio input device (dictation source), such as one USB microphone and one PowerMic II or PowerMic III, then the *Speech Node* needs to carry out a separate ACO process for each of those devices

After a user has dictated his or her first 5 hours of total audio, if the user has an accent, you might want to schedule the user for more frequent ACOs that you would schedule for users who do not have an accent.

Running a maintenance ACO process periodically keeps the user's master user profile synchronized with his or her voice.

Factors in how often to run LMO processes

Since an LMO process runs quickly, not tying up network or *Speech Node* resources, you can schedule an LMO for every user every day. Guidelines state at minimum to run an LMO process:

- Every day (or night) for users whose speech *Dragon* has difficulty recognizing or users who have **ever** had their profile becomes corrupted.
- At least every week (but preferably every day) for users new to dictating (who do not have a previously existing user profile or who have not yet dictated 5 full hours of audio).
- At least every month for users with more experience dictating (who already have a user profile and have already dictated a minimum of 5 hours of audio) and who have **never** had profile corruption issues.

Some of the factors to consider when deciding how often to carry out an LMO process are:

- How much dictation a user or group of users generates
- How often you add new words to vocabularies
- How often you would like the Nuance Management Server to check for corrupted user profiles, because it checks for corruption every time it carries out an LMO process on the profile

Factors that have no effect on regularly scheduled processes

After you approve medication and command updates, you do not take any other action. The Nuance Management Server automatically schedules processes that integrate these updates into the appropriate models and master user profiles.

Determining number of Speech Nodes required for optimizations

Another aspect of managing speech recognition is the process of integrating new speech data into the acoustic and language models in the user profiles.

The component of Dragon Medical Network Edition that updates and maintains acoustic and language models is the *Nuance Management Server*. DM Network Edition includes the *Nuance Management Server*. The *Nuance Management Server* always resides within your network, because it needs access to the Master User Profiles. The processes that integrate new information into the acoustic and language models for a particular healthcare user are called *acoustic model optimization* and *language model optimization*.

Before you try to determine how to include the *Nuance Management Server* in your system, you should determine how much space is required to carry out these optimizations, taking into account several factors.

Determining number of Speech Nodes required to process acoustic and language models

Once you know the total number of dictating users, you can then estimate how many *Profile Optimizer Speech Nodes* your network requires to regularly integrate and optimize the information *Dragon* learns about each user's speech patterns.

Each individual *Profile Optimizer Speech Node* spends about 20 hours (1200 minutes) in a given day carrying out the process of learning how each speaker speaks and adapting an acoustic model for each dictating user's voice based first on initial training, then later on actual day-to-day dictation. The acoustic model for a speaker needs to frequently incorporate corrections to speech recognition results. This process is called acoustic model optimization (ACO) the process is more or less continuous, depending on how many users you have. The number of minutes varies because the person's speech patterns vary in ways that are not predictable. And the greater the speed of the processor that each *Profile Optimizer Speech Node* runs on, the greater the amount of speech it optimizes in a day.

The *Profile Optimizer Speech Node* spends the remaining time in a day interacting with the *NuanceLink* and *NMS Server* to integrate new words into each professional's vocabulary, effectively modifying the language model for that user. The language model also incorporates statistical information about words and phrases most likely to occur in the context of the particular user's speech/-composition (writing) style and combines that information with known data about speakers of that user's language with the user's accent, if an accent is applicable. This process is called language model optimization (LMO).

The bulk of the *Profile Optimizer Speech Node*'s time is spent carrying out ACO. ACO essentially requires an hour of adaptation for an hour of dictation. So estimating the average amount of dictation per user gives you a good starting point to determining how many *Profile Optimizer Speech Nodes* you need.

The way that acoustic model optimization works is that although each *Profile Optimizer Speech Node* carries out 1200 minutes per day of adaptation, not all minutes of dictation are considered equal.

When a user has logged in for the first time and has been speaking into the microphone, after the user finishes dictating and logs out, the *Dragon Client* sends data about the user's speech to the corresponding master user profile on the network. If that user has spoken for at least 15 minutes, the *Profile Optimizer Speech Node* then carries out an ACO process on the acoustic model for that user

and updates that master user profile accordingly. Later, when the user logs on to dictate again, the *Dragon Client* retrieves the optimized master user profile from the network and places it in its local cache. The user should then see improved recognition.

After the user again dictates, then logs out of the session, the client again sends data about the user's speech to the master user profile on the network. The user may require more than one session to accumulate enough dictation, but when he or she has dictated at least 50 more minutes, the *Profile Optimizer Speech Node* automatically carries out another ACO process. (These transfers of data and ACO processes occur on the *Speech Node*, without the user even knowing about them.)

The first week that the user dictates, the *Profile Optimizer Speech Node* carries out an ACO process. The same occurs after another week. ACO processes become progressively further apart the longer a user dictates, until the user's acoustic model reaches maintenance mode, where the *Profile Optimizer Speech Node* carries out an ACO only as often as you schedule it and only if the user has dictated at least 5 hours of new speech during that period. For instance, you might schedule an ACO monthly; if the *Profile Optimizer Speech Node* does not find enough speech data for that user during that month, the node terminates the process and waits until the next ACO that is scheduled.

The table that follows estimates how many users a single node can process ACOs for various levels of dictation.

Number of minutes required for adaptation processes

Days of user Dictation	Stage of user "Training"	Minutes of Dictation per user Before ACO	Number of users One Speech Node Supports	Time Required to Tune Acoustic Model for Number of users with One Speech Node
First few days	First few days	15 minutes	_ 80 users	80 users X 15 minutes each = 1200 minutes in one day
End of 1st Week	Early dictation	50 minutes	_ 25 users	25 users X 50 minutes each = 1200 minutes per day
End of 2nd Week	More dictation	100 minutes	_ 12 users	12 users X 100 minutes each = 1200 minutes per day
End of Month		200 minutes	_ 10 users	10 users X 200 minutes each = 1200 minutes
Upgraded and in Maintenance mode	Upgrade training and ongoing dictation	5 hours (300 minutes)	250 users x 30% dictating at one time = 75 users	40 users X 30 minutes per day = processes acoustic models at 1200 minutes per day; two nodes can share ACO load during upgrade process

If a user uses more than one audio input device (dictation source), such as one USB microphone and one PowerMic II or PowerMic III, then the *Speech Node* needs to carry out a separate ACO process for each of those devices, so an extra audio input device is *almost* the same as another user. *Almost* the same, but not exactly the same, because a user who divides time between two devices probably spends less time dictating with each single device than another user spends dictating with only one device.

A single *Speech Node* can handle ACO and LMO for only 40 to 80 users if they are all just dictating for their first few weeks, but that same node can handle 250 users who have been dictating long enough to require an ACO be scheduled only once a month.

Probably not all of your users have the same level of experience with the speech recognition system, so you might have to make a judgment call. To help you make that call, you can enter your information into the table below and develop a grid that illustrates the number of *Speech Nodes* your installation should require.

Use the next table to estimate how many users a single node can process ACOs for at various levels of dictation. If you add up the number of users at each stage by the number minutes before an ACO occurs, you can determine how many total minutes of ACO processing you need. However, not all users are going to be dictating at the same time. For instance, if you have 240 users, and only about 30% of them (80) are dictating at one time, you need to multiply the total number of minutes by 30% before dividing it by the 1200 minutes in a day a single node can carry out ACO processes. A percentage of 30% is a good estimate of how many users dictate during a given period of time (such as a hospital shift); however, you can adjust the percentage if you believe as many as 50% or as few as 20% of your users dictate at any one time.

Number of Speech Nodes required in your installation

Time user Has Dictated	Approx. Number of users at This Stage During This Month	Minutes of Dictation Before ACO	Multiplier of Min/Day	TOTAL minutes
2 Days		15 min	X 15	
1 Week		30 min	X 30	
2 Weeks		1 hr	X 60	
Month		2 hrs	X 120	
Being Upgraded		60 min	X 60	
More Than		5 hrs	X 30	

Time user Has Dictated	Approx. Number of users at This Stage During This Month	Minutes of Dictation Before ACO	Multiplier of Min/Day	TOTAL minutes
One Month (Maintenance)				
TOTALS				_____ x _____ % /1200 = _____ Nodes

Recommendations about number of users

A good rule of thumb for how many users a single *Speech Node* on a dual-quad core machine can process ACOs for is 250 users who are in maintenance mode per pair of processor cores beyond the first core. (The first core is reserved for all other server processes of the DM Network Edition network.)

If you have a quad core machine dedicated to *Speech Nodes*, you can create two virtual machines, two cores each, and run a single *Speech Node* on each virtual machine.

The following table delineates how many *Speech Nodes* you should need as your network matures and your *Profile Optimizer Speech Nodes* run ACOs mostly in maintenance mode.

Number of Speech Nodes Recommended to Optimize Specific Numbers of Users at Particular ACO Intervals

Number of users	Interval between Acoustic Model Optimizations (ACOs) in Months					
	1 mth recommended	2 mths	3 mths	4 mths	5 mths	6 mths minimum
0 – 240	1	1	1	1	1	1
241 – 480	2	1	1	1	1	1
481 – 720	3	2	1	1	1	1
721-960	4	2	2	1	1	1
961 – 1200	5	3	2	2	1	1
1201 – 1440	6	3	2	2	2	1
1441 – 1680	7	3	3	2	2	2
1681 – 1920	8	3	3	2	2	2
1921 – 2160	9	4	3	3	2	2
2161 – 2400	10	4	4	3	2	2
2401 – 2640	11	5	4	3	3	2

That quad-core processor lets you install one virtual machine for *Speech Nodes* on the physical server or workstation. The number of users you can process with is from 481 to 720 (up to 240 users per node) if you are running an ACO on those users every month. However, if you run the ACO only once every two months, you can increase the number of users up to between 960 and 1200.

If you have the *Speech Nodes* on their own machine, separate from the servers, you can have four virtual machines on an eight core physical server, each running a single node, and all those *Speech Nodes* can together run ACOs every two months for from 2,161 to 2,400 users.

Nuance recommends that you run an ACO on every user once a month to maintain accurate speech recognition. The absolute minimum number of ACOs you should run for a single user is one every 6 months, but Nuance recommends more frequent tuning of the acoustic models.

Estimating storage requirements for the Nuance Management database database

About Nuance Management Server (NMS) and NMS SQL database

If you are running the NMS within your network, the NMS SQL Database is within your network. The *NMS SQL Database* stores information about the customer account for the organization as well as sites, users, groups, and licenses.

Most of this information makes a scant difference in the size of the *NMS* database. *NMS* generates the bulk of the data by taking actions such as auditing events, logging statuses, managing scheduled tasks, and storing messages and packages it receives from Nuance servers.

The 20 GB free space requirement for the *NMS* installation provides ample space for not only installation of the *NMS SQL Database*, the *NMS Server*, and the *Nuance Management Center*, but for all data generated at most installations with up to 500 dictating users.

You do not have to allocate any space on the *NMS Server* for data collection. When you choose to implement data collection (it is optional) to help Nuance improve Dragon's future speech recognition ability, the server does not collect the data; instead each *Dragon Client* on the network collects data while the user dictates, then sends the appropriate data from the local user cache profile to the computer hosting the master user profiles. The *NMS Server* retrieves and packages applicable data from the master user profiles, then sends it to Nuance every night. Because it is utilizing data

already stored in the master user profiles, the *NMS Server* does not require additional storage space for data collection.

About disk space for server installations and logs

In addition to considering how much space a database requires, you should also be planning for space for each server installation and for log files.

The NMS server requires 20 GB of free space.

In addition, each *Profile Optimizer Speech Node* generates two log files for each ACO and LMO process it carries out:

- Dragon log file—*Dragon Medical SDK Client* writes this file
- *Speech Node* service log file—*Profile Optimizer* services write these files

These files are generally between 3 and 5 KB each. The *Nuance Management Server (NMS)* stores them for 30 days. Although the log files could grow larger, they are stored in a zipped format and compress very well. If you run an optimization process for every user every day and have 500 users, then assume every log will be as large as 5 KB, all of those logs would still use only 2.5 MB of storage space and in 30 days that would become 75 MB of storage space. After 30 days, the *NMS* purges old log files, so the storage requirements do not grow substantially.

The *NMS Server* generates Windows communication foundation service log files for each user and stores up to the last 10 MB of them generated, a maximum of 50 log files for the entire server, before purging the oldest files and replacing them with newer ones as it generates them.

Storage space for all of these logs is relatively inconsequential; however, if you expand the number of users in your organization, it is always best to take into account storage requirements for all generated files.

Recommended network and switch settings

Nuance recommends particular network interface card settings and network speed for equipment in the Dragon Medical Network Edition network.

Network Interface Card (NIC) settings

Gigabit Cards: Gigabit cards should be set to automatic. The network switches and the cards plugged into them should have the same setting.

10 Mb/100 Mb Cards: Network link speed and duplex need to be set the same on all servers, workstations, hubs, switches, or other network equipment. If there is a mismatch in settings, or if the NICs are left set at **Auto Detect**, the end users of the system could see degradation in both performance and recognition. Nuance does not support the **Auto Detect** setting for 10/100 Mb cards.

Network speed

Nuance supports all network speeds/settings listed below when you set them consistently across the network:

- 100 Mbps/Full Duplex
- 10 Mbps/Full Duplex

Regardless of 100 or 10, matching the **Full Duplex** setting is required. The network switches and the cards plugged into them should have the same setting.

Using Network traffic switch for load balancing

If you use multiple *NMS Servers*, you can insert a network traffic switch, such as the one available through F5 and similar manufacturers, into your DM Network Edition network for balancing the load distribution among those servers. For details on the exact message that the network traffic switch can send to the *NMS Server* to ping it, refer to the DM Network Edition *Installation Guide* or the *DM Network Edition Administrator guide*.

Storage hardware requirements for master user profiles

For details on optional hardware you can deploy to store your master user profiles, refer to *System requirements for master user profile computers* on page 44.

Ports to open for clients, servers, and hardware firewalls

On the Dragon Medical Network Edition network, you must open particular ports to ensure the free flow of data between servers, clients, and applications. You must open all required ports on any hardware firewalls to ensure that firewalls do not block transmission of data.

Installing the product without upgrading from a previous version of the product

The NMS uses the standard port 443 to communicate with the Dragon client.

You should allow/forward the original non-standard ports, as well as the standard port 443. This should provide maximum flexibility until all your Nuance client applications that support the standard port 443.

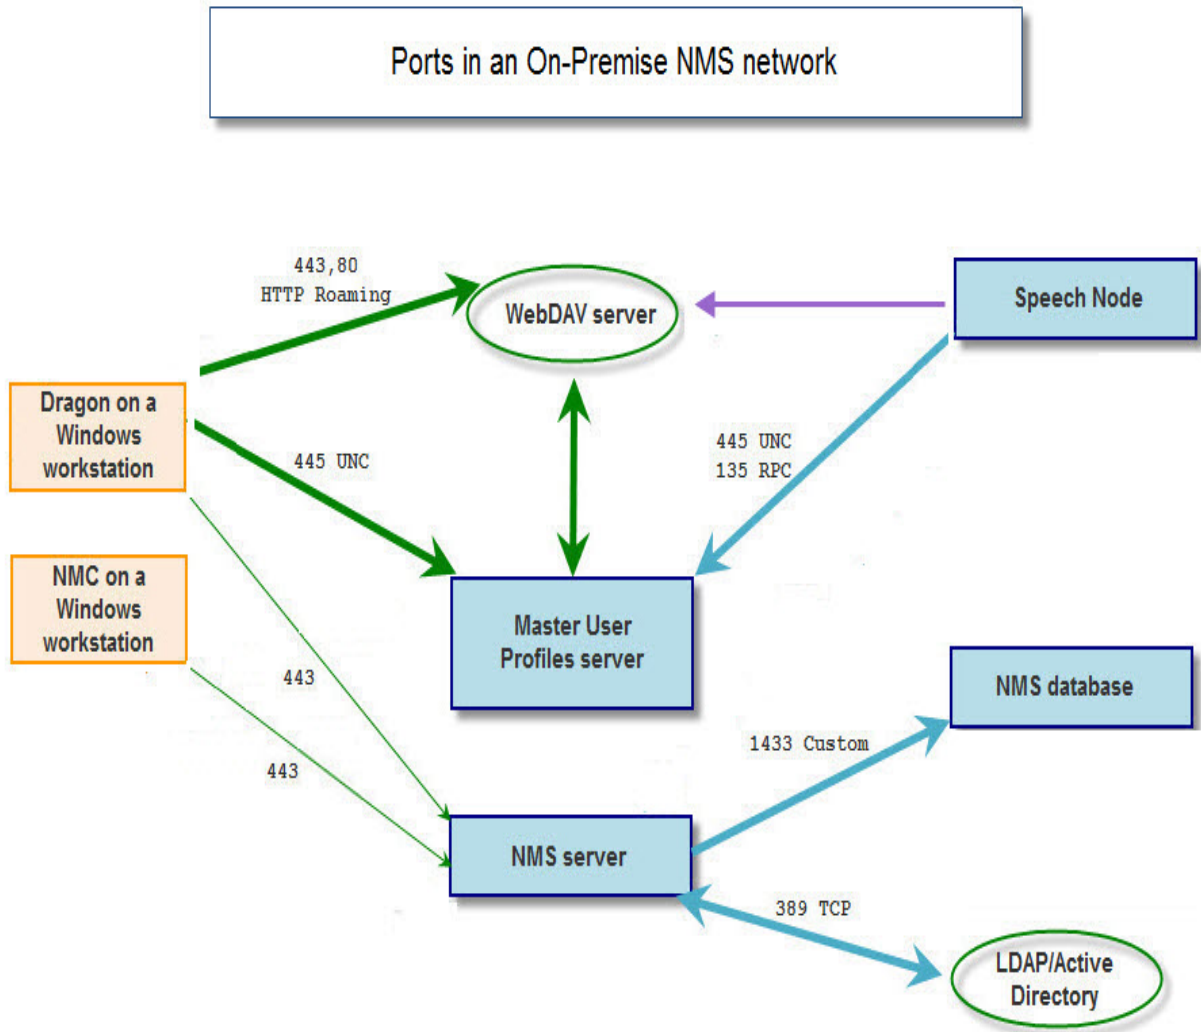
Upgrading from a previous version of the product

When you upgrade, the NMS does not automatically enable the standard port 443.

You can use the .config file to reconfigure the NMS to perform any forwarding only for port 443.

You should not modify your deployment of Nuance products until all of those products support the use of the standard port 443. At that time, you may modify the NMS configuration to no longer support the older ports.

Ports to open when using an on-premise NMS



You must open the following ports in the DM Network Edition network:

- Ports between *Dragon Client* workstations and the *Master User Profiles Server*
- Ports between workstations where you expect to run the *Nuance Management Center* and the *NMS Server*
- Ports between *NMS Server* and its database if that database is on a separate server
- Ports between each *Speech Node* machine and the *Master User Profiles Server*
- Ports between each *NMS Server* and the *Nuance Update* server.
- All ports mentioned above on all hardware firewalls protecting the network

Chapter 3: Configuring DM Network Edition networks that use physical servers

This chapter presents information about different configurations for DM Network Edition networks that primarily use physical servers.

Using physical servers in a DM Network Edition network	34
System requirements for small networks with physical servers	39
System requirements for medium networks with physical servers	40
System requirements for large networks with physical servers	42
System requirements for master user profile computers	44
Summary: System requirements for DM Network Edition networks that use physical servers	46
Summary: Configuration options for DM Network Edition networks that use physical servers ...	49

Using physical servers in a DM Network Edition network

Configuring a DM Network Edition network for up to 100 users

Use the following system requirements as guidelines when you set up a network for up to 100 users.


You can have more than one master user profiles directory if you have more than one site. For more information, refer to the DM Network Edition Administrator guide.

The *NMS Server* runs more efficiently on a quad-core server because it might need to run compression/encryption on up to 600 MB of data for a single user profile to support data collection. In addition, if you are storing them on the server, the master user profiles need to be on their own RAID array. Alternatively, you can place the *Master user profiles* directory on its own entirely separate machine, either a physical server or a workstation.

Configuring a DM Network Edition network for 100 to 1000 users

The most optimal distributed configuration installs the *NMS* on one physical server and installs the *Speech Nodes* on a separate physical server:

- First server software components:
 - *NMS Server* service with its Nuance Management Center
 - *NMS SQL Database*
 - Master user profiles directory (if not stored elsewhere)
- Second server software components:
 - Up to 4 *Profile Optimizer Speech Nodes*, each on its own virtual machine.



Configuring a DM Network Edition network for more than 1000 users

If you have a larger organization with more than 1,000 dictating users, you need to scale the network by having at least three physical servers running multiple *NMS Servers* and multiple *Speech Nodes* on virtual machines.

Each of the physical servers in this configuration must be an eight-core server:

- At least one dual-core server with these software components. If you are not using the NMS on-premise, you don't need this server:
 - One NMS Server for every 1,000 users
 - A *Nuance Management Center* for each *NMS Server*
- A single Database server with:
 - *NMS SQL Database*
 - Master User Profiles directory (if not stored elsewhere)
- At least one eight-core server with these components:
 - Up to 4 *Profile Optimizer Speech Nodes*, one for every 250 users

Note:

You can have more than one sub-directory inside the master user profiles directory, one for each site. For more information about setting up sites, refer to the DM Network Edition Administrator guide.

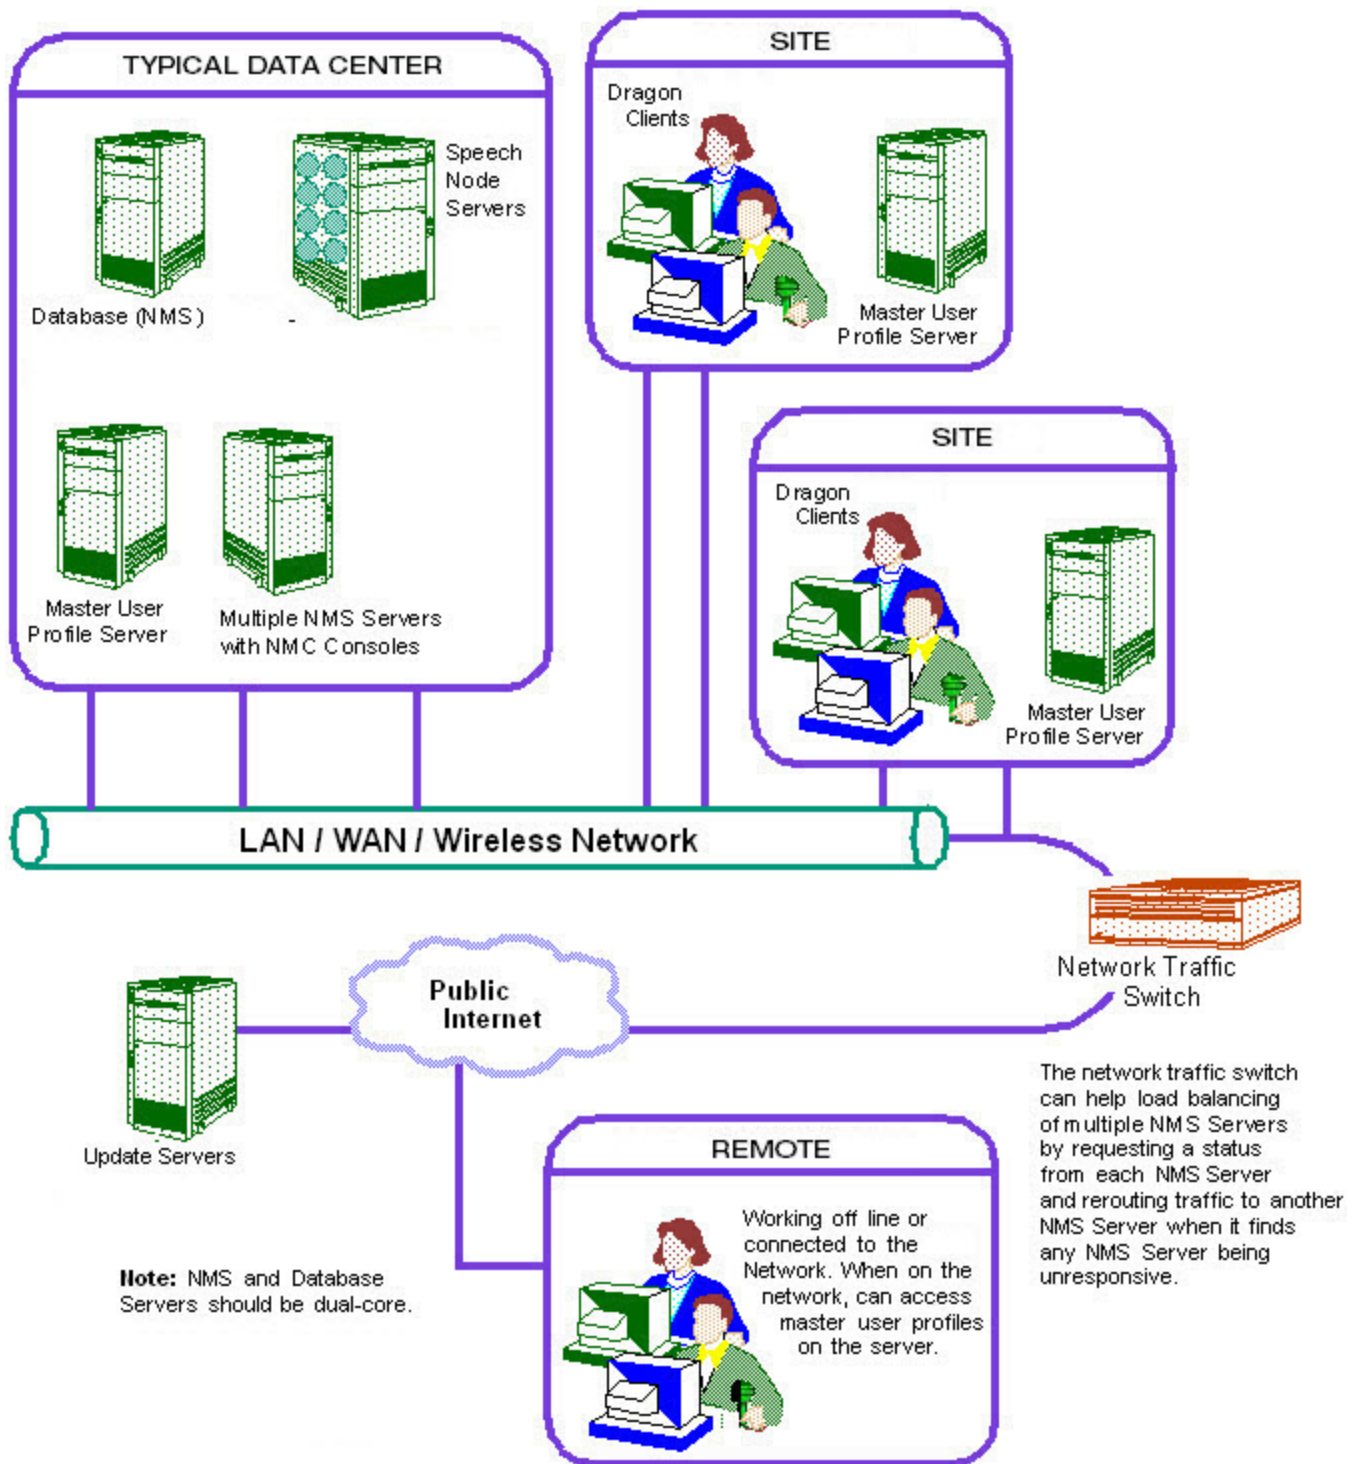
In a large distributed configuration, you can have multiple separate distributed *Master User Profile Servers* so that clients can have quick local access to the user profiles. In a situation like this, a connection between the client machine and the local *Master User Profile Server* can be much slower than the network connection between distributed sites; for instance, 10 Mbps might be an adequate speed for the local area network, whereas 100 Mbps might be required for the wide area network to other sites.

Note: Speech node servers need two cores for every Speech Node, for example quad-core for two Speech Nodes.

The *NMS Server* must be dual-core, because the *NMS Server* takes advantage of the extra core to provide management of your Dragon Medical Network Edition network. The *Speech Node* server hardware must have two cores for each *Speech Node*, so it is best to have an eight-core server for every four nodes.

The next illustration shows the various servers and clients in a DM Network Edition network configuration for 4,000 users. Note that this configuration contains a large data center with a hospital site, a clinic site, and a network traffic switch to manage *NMS Server* load balancing. The network traffic switch (available through F5 and similar manufacturers) can send a request to each *NMS Server* to query its status, tag a server as down if it does not respond, and reroute the traffic to other *NMS Servers*.

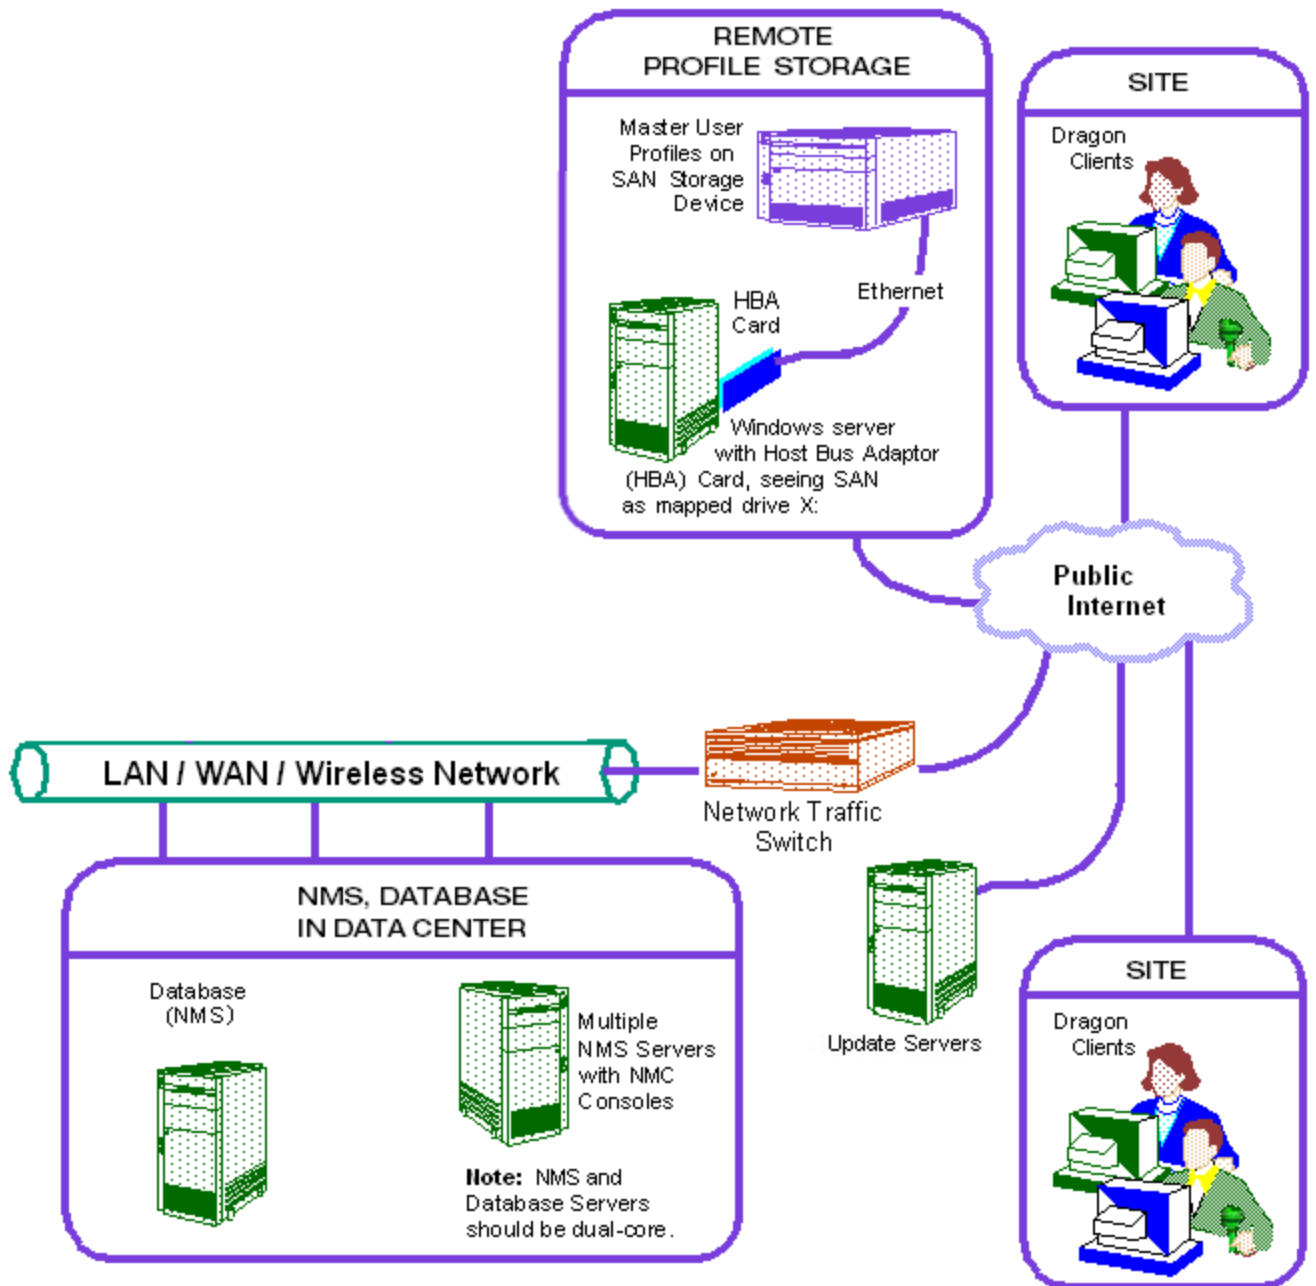
Network Configuration using large data center, for approx. 4,000 dictating users



Notice that a *Master User Profile Server* local to the client provides rapid access to user profiles.

The next illustration shows the master user profiles stored on a remote SAN storage device.

Remote Master User Profiles on SAN storage device



System requirements for small networks with physical servers

Use the following system requirements as guidelines when you set up a network for up to 100 users.

■ Single Quad-Core Physical Server

- **One Quad-Core Server** for NMS Server, Nuance Management Center, NMS SQL database, and One Speech Node.
- Master User Profiles Directory on independent machine (server not required)

■ Hardware and Software

- **Processor:** Quad-Core 2 GHz CPU
- **Operating System:**
 - Windows Server 2008 or with SP2 (optional)
 - Windows Server 2008 R2
 - Windows Server 2012
 - Windows Server 2012 R2
- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5

■ Minimum RAM: 4 GB recommended

■ Core Application Disk Storage: 40 GB for the NMS server

■ Maximum Number of Speech Nodes:

- If using only 1 physical server, it needs to be a Virtual Server Host. NMS may share a Virtual Machine. The Speech Node must be on a separate Virtual Machine
- One *Profile Optimizer Speech Node*
- *Speech Node* uses two full cores of the CPU and must be on a Virtual Machine if deployed on a single server.

■ Minimum and maximum number of users: 10 to 100 dictating users

■ Master User Profile Storage:

- Master user profiles directory on its own RAID array.
- Master user profiles on their own totally separate server or workstation

System requirements for medium networks with physical servers

Use the following system requirements as guidelines when you set up a network for 100 to 1000 users.

Two Physical Servers, One Quad-Core and One Eight-Core

- **First Server (Quad-Core):** NMS Server, NMC Center, NMS SQL database,.
- **Second Server (Eight-Core):** Up to 4 *Profile Optimizer Speech Nodes* (one for each two cores)
- Master User Profiles Directory on a server or on a separate machine (server not required, but recommended)

First server requirements

- **Processor:** Dual-Core 2 GHz CPU
- **Operating System:**
 - Windows Server 2008 with SP2 (optional) or R2
 - Windows Server 2012
 - Windows Server 2012 R2
- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5
- **Minimum RAM:** 4 GB RAM for NMS
- **Approximate Disk Storage:** 20 GB for the NMS Server
- **Master User Profile Storage:**
 - Master User Profiles Directory on a server, a separate machine, or on its own RAID array (server not required, but recommended)
- **Second server requirements:**
 - **Processor:** Eight-Core 2 GHz CPU
 - **Operating System:**
 - Windows Server 2008 or with R2. On all machines where *Speech Nodes* will run on Windows Server 2008 R2 (without Service Pack 1), download and apply the following patch:
http://hotfixv4.microsoft.com/Windows%207/Windows%20Server2008%20R2%20SP1/sp2/Fix387409/7600/free/441462_intl_x64_zip.exe

Please see the following Microsoft Support articles for information about

why this hotfix is necessary for reliable operation of a Speech Node on Windows Server 2008 R2:

- <http://support.microsoft.com/kb/2624677>
- <http://support.microsoft.com/kb/2687753>

- Windows Server 2012
- Windows Server 2012 R2

- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5

- **Minimum RAM:** 8 GB RAM

- **Minimum Disk Space:** 20 GB hard disk space for each Speech Node

- **Maximum Number of Speech Nodes:** Up to 4 *Profile Optimizer Speech Nodes* per eight-core server, one node for each two cores on the server

- **Minimum and maximum number of users:** 101 to 1,000 dictating users

System requirements for large networks with physical servers

Use the following system requirements as guidelines when you set up a network for more than 1000 users.

One NMS Server for every 1,000 users

- **One Database Server** for NMS SQL database (if used) and Master User Profiles Directory
- **One Eight-Core Server** for every 4 *Profile Optimizer Speech Nodes*

First server requirements

- **Processor:** Dual-Core 2 GHz CPU
- **Operating System:**
 - Windows Server 2008 with SP 2 (optional) or with R2
 - Windows Server 2012
 - Windows Server 2012 R2
- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5
- **Minimum RAM:** 4 GB RAM
- **Approximate Disk Storage:** 20 GB for each *NMS Server* instance
- **Database server requirements:**
 - **Processor:** Dual-Core 2 GHz CPU
 - **Operating System:**
 - Windows Server 2008 with SP 2 (optional) or with R2
 - Windows Server 2012
 - Windows Server 2012 R2
 - **Minimum RAM:** 4 GB RAM
 - **Master User Profile Storage:**
 - Master user profiles directory on its own RAID array
- **Eight-Core Server requirements:**
 - **Processor:** Eight-Core 2 GHz CPU
 - **Operating System:**

Windows Server 2008 or with R2. On all machines where *Speech Nodes* will run on Windows Server 2008 R2 (without Service Pack 1), download and apply the

following patch:

http://hotfixv4.microsoft.com/Windows%207/Windows%20Server2008%20R2%20SP1/sp2/Fix387409/7600/free/441462_intl_x64_zip.exe

Please see the following Microsoft Support articles for information about why this hotfix is necessary for reliable operation of a Speech Node on Windows Server 2008 R2:

- <http://support.microsoft.com/kb/2624677>
- <http://support.microsoft.com/kb/2687753>
- Windows Server 2012
- Windows Server 2012 R2
- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5
- **Minimum RAM:** 8 GB RAM
- **Approximate Disk Storage:** 20 GB for each *Profile Optimizer Speech Node*
- **Minimum and maximum number of users:** Up to 1,000 dictating users per eight-core server.
- **Maximum number of speech nodes:** Four *Profile Optimizer Speech Nodes* per eight-core server, one node for each two cores on the server

System requirements for master user profile computers

Any time after you have installed the *NMS Server* and started the *NMS Server* service, you can set up your *Master user profile storage* machine.

Hardware and software requirements

Hardware requirements

Processor:

Intel® Pentium 4® or later or AMD Athlon 64 1 GHz processor or later.

Recommended: 2.4 GHz processor or greater.

Faster processors produce faster performance.

RAM:

4 GB RAM for Windows Server 2008, Windows Server 2012, or Windows Server 2012 R2.

Cache:

512 KB minimum L2 Cache. *Recommended:* 1 MB L2 Cache.

Software requirements

Operating system options for host machine:

You can store your master user profiles on a machine running one of these operating systems:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

Other software required:

The Windows-based machine where you store the master user profiles must have .NET Framework 4.5.

Optional hardware

Storing Master User Profiles on RAID array

If you want or need to store the profiles on a RAID array, that array must be connected to a Windows-based machine with .NET Framework 4.5.

Storing Master User Profiles remotely on a SAN storage device

You can store your master user profiles at a remote location for access from several sites by storing them on a SAN storage device. To set up this kind of topology, when you store the master user profiles on a SAN storage device, the SAN device must be connected to a Dragon Medical Network Edition server or another network server running Windows 2008 or 2012 Server with .NET Framework 4.5 over a 10/100/1000 Mbps Ethernet using a Host Bus Adapter (HBA) card. The server that connects to the SAN must access the SAN storage device as a mapped drive (such as X:).

Options for storing user profiles on machines running server components

You can store your master user profiles on either a server machine or a workstation NMS Server, and the NMS SQL Database.

Optional software

In a typical installation, you would store your master user profiles on a domain (network) machine and access them through either a mapped disk drive or a UNC. Other alternatives include storing your user profiles on either a web server or a secure web server.

Storing your user profiles on a web server or secure (SSL) web server

You can store your master user profiles on a web server or a secure web server.

- *Dragon Client* workstations would access the profiles on a web server through a URL starting with **http://**.
- *Dragon Client* workstations would access the profiles on a secure (SSL) web server through a URL starting with **https://**.
- *NMS Server* and *NMC Center* would access the profiles through a UNC path that points to the user profile location.

For more information on compatible web servers and installing the software required to set them up, refer to the DM Network Edition *Installation Guide*.

Summary: System requirements for DM Network Edition networks that use physical servers

System requirements for databases, the NMS Server, and the NMC Center

	NMS SQL Database	NMS server
Computer Type	Server	Server
Processor	Xenon 3050 or equivalent	Xenon 3050 or equivalent
CPU Speed	2 GHz	2 GHz
Gigabit Network Cards	Gigabit cards and switches/cards plugged into them = Automatic	
10 Mb/100 Mb Cards	Same network link speed & duplex settings on all servers, workstations, hubs, switches, and cards.	
Network Speed	100 Mbps/Full Duplex or 10 Mbps/Full Duplex	
Minimum free hard drive space	20 GB	20 GB
Minimum RAM	4 GB	4 GB
Minimum L2 cache	2 MB	2 MB
Other Hardware items	---	---
Internet Information Services (IIS)	---	Required
In Windows Domain	Required	Required
Windows Installer	---	---

	NMS SQL Database	NMS server
3.1 or later		
Dragon SDK Client	---	---
Internet Browser	---	Internet Explorer 8, 9, 10, and 11, Chrome, Firefox, or Safari.

System requirements for Speech nodes, Master User Profiles Host, and the Dragon client

	Master User Profiles Host	Speech nodes	Dragon client
Computer Type	Workstation or Server	Workstation or Server	Workstation
Processor	At least Pentium IV	Two cores per node	Pentium IV or AMD Athlon 64
CPU Speed	2 GHz	2 GHz	2.4 GHz (AMD 1 GHz)
Gigabit Network Cards	---	---	---
10 Mb/100 Mb Cards	---	---	---
Network Speed	---	---	---
Minimum free hard drive space	Based on calculation	20 GB	5 GB
Minimum RAM	4 GB	2 GB per node, up to 8 GB on an eight-core server	2 GB RAM for: -Windows XP -Windows Vista 4 GB RAM for: -Windows 7 -Windows 8 and 8.1 -Windows 10 -Windows Server 2008 64-bit -Windows Server 2012 64-bit
Minimum L2 cache		2 MB	512 KB; Recommended: 1 MB L2
Other		---	Sound card recording

	Master User Profiles Host	Speech nodes	Dragon client
Hardware items			at 16 bit 11 KHz and speakers for playback; Microphone; DVD reader
Windows Operating System	Windows XP* SP2 or SP3, Windows Vista, or Windows 7 workstation; Windows Server 2008 , Windows Server 2012, and 2012 R2; RAID Array may be required (see note below)	Windows Server 2008 and 2008 R2 (64-bit), Windows Server 2012, and 2012 R2	32-bit and 64-bit operating systems detailed below table.
SQL Server		---	---
.NET Framework		Version 4.5	Version 4.5
Internet Information Services (IIS)		---	----
In Windows Domain		Not required	Required
Windows Installer 3.1 or later		Required	---
Dragon SDK Client		Required	---
Internet Browser		---	

Notes:

Summary: Configuration options for DM Network Edition networks that use physical servers

	Size and type of network configuration		
	Small: One server	Medium: Two servers	Large: Three or more servers
Number of users	10 - 100	101 - 1,000	1,000 or more
Minimum Physical Servers	1 Server	2 Servers	Servers: One Dual-Core NMS Server & one Eight-Core Speech Node Server for every 1,000 users (one Speech Node server allows for 4 Speech Nodes); One Database Server
NMS Server			
Software	-NMS Server -NMS database -Master User Profiles Folder -1 Speech Node	-NMS Server -NMS database -Master User Profiles Folder	-NMS Server
Hardware	-Quad-Core with minimum 4 GB RAM -Hard Disk = 40 GB + 1 GB per user	-Dual-Core with minimum 8 GB RAM -Hard Disk = 40 GB + 1 GB per user	-Dual-Core with minimum 4 GB RAM -Hard Disk = 40 GB + 1 GB per user
Database Server			
Software	---	---	-NMS database -Master User Profiles Folder
Hardware	---	---	-Dual-Core with minimum 4 GB RAM -Hard Disk = 20 GB per node + 1 GB per user
Optional Multiple Distributed Master User Profile Servers, one for each site (for large systems with distributed sites)			
Software	---	---	Master User Profiles Folder
Hardware	---	---	Intel® Pentium 4® or later or AMD Athlon 64 1 GHz processor or later with from 1 GB RAM to 4 GB RAM, depending on Operating System. Refer to <i>System requirements for master user profile computers</i> on

	Size and type of network configuration		
	Small: One server	Medium: Two servers	Large: Three or more servers
			page 44. -Hard Disk = 20 GB + 1 GB per user

Chapter 4: Configuring DM Network Edition networks that use virtual servers

This chapter presents information about different configurations for Dragon Medical Network Edition networks that primarily use virtual servers.

Using virtual servers in a DM Network Edition network	52
Special Requirements when running Dragon in a virtual environment	53
System requirements for small networks with virtual servers	54
System requirements for medium networks with virtual servers	55
System requirements for large networks with virtual servers	58

Using virtual servers in a DM Network Edition network

To use virtual machines in your Dragon Medical Network Edition network configuration, you should be deploying at least a medium-sized configuration—one with the Profile Optimizer Speech Nodes on a separate physical server or workstation.

Using Virtual Machines for the NMS Server, and databases

You can choose to generate medium or larger configurations using virtual machines for the NMS Server, and databases: Hardware, operating system, and software requirements remain the same on virtual machines.

Using Virtual Machines for Profile Optimizer Speech Nodes

You can have each *Profile Optimizer Speech Node* on its own virtual machine as long as the physical workstation or server for that virtual machine has two full cores for each *Speech Node*. On a dual-core machine, you can have one *Speech Node*; on a quad-core machine, you can have two *Speech Nodes* each on its own virtual machine.

Tested virtual machine desktop configurations for Dragon and VMWare

Operating system	Hardware Configuration
Windows 7 Professional 64-bit	1 vCPU <ul style="list-style-type: none">• 4 GB memory• 1 Ethernet card (e1000 driver)• 32 GB storage (vmdk only)

Special Requirements when running Dragon in a virtual environment

Make sure the system requirements for Dragon are met. See *System requirements for the Dragon Client* on page 62.

Only use BestMatch IV speech models.

The Dragon (natspeak.exe) process requires 1.1GB of RAM

Each active speaker requires a single (non-hyper-threaded) core

System requirements for small networks with virtual servers

Use the following system requirements as guidelines when you set up a network for up to 100 users.

- **Single Quad-Core physical or virtual server**

- **One Quad-Core Server** for NMS Server, Nuance Management Center, NMS SQL database, and One Speech Node.

- **Hardware and Software**

For this configuration, you need a more robust server machine than you would require for an installation distributed over several machines.

- **Processor:** Quad-Core 2 GHz CPU
- **Operating System:**
 - Windows Server 2008 or with SP2 (optional), 32- or 64-bit version
 - Windows Server 2012
 - Windows Server 2012 R2
- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5
- **Minimum RAM:** 4 GB is recommended
- **Core Application Disk Storage:**
40 GB for the Operating system and the NMS Server
- **Maximum Number of Speech Nodes:** One Profile Optimizer Speech Node that uses two full cores of the CPU. Must be on a Virtual Machine if deployed on a single server.
- **Minimum and Maximum Number of Users:** 10 to 100 dictating users
- **Master User Profile Storage:** 2.1 GB per user/dictation device is required for the master user profiles. Storage can be on the server, on another file store, a storage array, or on a storage area network (SAN).

System requirements for medium networks with virtual servers

Use the following system requirements as guidelines when you set up a network for 100 to 1000 users.

- **First Server (Dual vCPU):** NMS Server, Nuance Management Center, NMS SQL database
- **Additional Servers (1 per 250 Users) :** Speech Node
- Master User Profiles Directory on first server or on a separate Windows server or Distributed File System (DFS)
- **First server requirements:**
 - **Processor:** 2 vCPU's
 - **Operating System:**
 - Windows Server 2008 or R2, 32 bit or 64 bit
 - Windows Server 2012
 - Windows Server 2012 R2
 - **Microsoft .NET**
 - Microsoft .NET Framework 4.5
 - **Minimum RAM:** 4 GB RAM, 8 GB if SQL installed on the server
 - **Core Application Disk Storage:** 40 GB includes 14 GB for the Operating system
- **Second server (Speech Nodes 1 per 250 Users) requirements**
 - **Processor:** 2 vCPU
 - **Operating System:**
 - Windows 10, 32-bit or 64-bit
 - Windows 8.1, 32-bit or 64-bit
 - Windows 8, 32-bit or 64-bit
 - Windows 7, 32-bit or 64
 - Windows XP, 32-bit or 64
 - Windows Server 2012
 - Windows Server 2012 R2
 - Windows Server 2008, 32-bit or 64-bit

Windows Server 2008 R2: On all machines where Speech Nodes run on Windows Server 2008 R2 (with or without Service Pack 1), download and apply the following patch:

<http://hotfixv4.microsoft.com/Windows%207/Windows%20Server2008>

http://support.microsoft.com/fix392120/7600/free/445413_intl_x64_zip.exe

Please see the following Microsoft Support articles for information about why this hot fix is necessary for reliable operation of a Speech Node on Windows Server 2008 R2:

<http://support.microsoft.com/kb/2687753>

<http://support.microsoft.com/kb/2624677>

<http://support.microsoft.com/kb/2584874>

- **Microsoft .NET** : Microsoft .NET Framework 4.5

- **Minimum RAM:**

1 GB RAM for Windows XP

2 GB RAM for:

- Windows XP
- Windows Vista SP2
- Windows 7
- Windows 8
- Windows 10
- Windows Server 2008
- Windows Server 2012
- Windows Server 2012 R2

- **Hard Disk Space:**

20 GB for Windows XP

30 GB for:

- Windows Vista SP2
- Windows 7
- Windows 8
- Windows 10
- Windows Server 2008
- Windows Server 2012
- Windows Server 2012 R2

- **Master User Profile Storage:** 2.1 GB per user/dictation device is required for the master user profiles. Storage can be on the first server, a storage array, or on a storage area network

(SAN). For a SAN configuration, Fiber Channel connectivity to the server running the Profile Optimizer Speech Nodes is preferred to reduce network usage.

System requirements for large networks with virtual servers

Use the following system requirements as guidelines when you set up a network for more than 1000 users.

- **One NMS Server** for every 1,000 users
- **One Database Server** for the NMS SQL database and a customer SQL cluster
- **One Speech Node:** one node for every 250 users
- **First server: NMS Server hardware requirements and limitations**
 - **Processor:** 2 vCPU
 - **Operating System:**
 - Windows Server 2008, R2, or SP2 (optional), 32- or 64-bit version
 - Windows Server 2012
 - Windows Server 2012 R2
 - **Microsoft .NET:**
 - Microsoft .NET Framework 4.5
 - **Minimum RAM:** 4 GB RAM
 - **Approximate Disk Storage:** 40 GB
- **Database server: Database server hardware requirements and limitations**
 - **Processor:** 2 vCPU or Existing SQL Cluster
 - **Operating System:**
 - Windows Server 2008, R2, or SP2 (optional), 32- or 64-bit version
 - Windows Server 2012
 - Windows Server 2012 R2
 - **Minimum RAM:** 4 GB RAM, 8 GB is recommended
 - **Approximate Disk Storage:** 40 GB
- **Speech Nodes: (Speech Nodes 1 per 250 Users)**
 - **Processor:** 2 vCPU
 - **Operating System:**
 - Windows 10, 32-bit or 64-bit
 - Windows 8.1, 32-bit or 64-bit
 - Windows 8, 32-bit or 64-bit

- Windows 7, 32-bit or 64-bit
- Windows Vista, 32-bit or 64-bit
- Windows XP, 32-bit or 64-bit
- Windows Server 2008 or R2, 32-bit or 64-bit. On all machines where *Speech Nodes* will run on Windows Server 2008 R2 (without Service Pack 1), download and apply the following patch:
http://hotfixv4.microsoft.com/Windows%207/Windows%20Server2008%20R2%20SP1/sp2/Fix387409/7600/free/441462_intl_x64_zip.exe

Please see the following Microsoft Support articles for information about why this hotfix is necessary for reliable operation of a Speech Node on Windows Server 2008 R2:

- <http://support.microsoft.com/kb/2624677>
 - <http://support.microsoft.com/kb/2687753>
 - Windows Server 2012
 - Windows Server 2012 R2
- **Microsoft .NET:**
 - Microsoft .NET Framework 4.5
 - **Minimum RAM:**
 - 1 GB RAM for Windows XP

 - 2 GB RAM for:
 - Windows Vista
 - Windows 7
 - Windows 8
 - Windows 10
 - Windows Server 2008
 - Windows Server 2012
 - Windows Server 2012 R2

- **Hard Disk Space:**
 - 20 GB for Windows XP

30 GB for:

- Windows Vista
 - Windows 7
 - Windows 8
 - Windows 10
 - Windows Server 2008
 - Windows Server 2012
 - Windows Server 2012 R2
- **Minimum and Maximum Number of Users:** Over 1000 dictating users
 - **Master User Profile Storage:** 2.1 GB per user/dictation device is required for the master user profiles. Storage can be on the first server, a storage array, or on a storage area network (SAN). For a SAN configuration, Fiber Channel connectivity to the server running the Profile Optimizer Speech Nodes is preferred to reduce network usage.

Chapter 5: System requirements for the Dragon client

This chapter presents the system requirements for the DM Network Edition Dragon client. Use the information in this chapter as guidelines when you set up your DM Network Edition network.

System requirements for the Dragon Client	62
--	-----------

System requirements for the Dragon Client

During the installation process, the Dragon client install wizard checks that your system meets the following requirements. If they are not met, Dragon is not installed.

Hardware Recommendations

Dragon Client Hardware Recommendations

CPU: 2.4 GHz Intel Dual Core or equivalent AMD processor. (IMPORTANT: SSE2 instruction set required)

Processor Cache: 2 MB

DVD-ROM: drive required for installation

Sound Card: Creative® Labs Sound Blaster® 16 or equivalent sound card supporting 16-bit recording.

Free hard disk space: 5 GB required, 8 GB recommended

RAM:

2 GB for:

- Microsoft® Windows XP®
- Microsoft® Windows Vista®

4 GB for:

- Microsoft® Windows® 10, 32-bit and 64-bit
- Microsoft® Windows® 8 and 8.1, 32-bit and 64-bit
- Microsoft® Windows® 7, 32-bit and 64-bit
- Microsoft® Windows Vista®, 32-bit and 64-bit
- Windows Server 2012, 64-bit R2.
- Windows Server 2008, 64-bit, Windows Server 2012 64-bit

Microphone: Nuance-approved microphone (included in purchase)

For details on Bluetooth microphones, recorders, Tablet PCs, and other hardware, please go to <http://support.nuance.com/compatibility/>.

You can also use an iOS or Android device as a microphone using the [Dragon Remote Microphone app](#).

Bluetooth (Optional): For Bluetooth wireless microphone support, visit <http://support.nuance.com/compatibility/>

Web Connection: A web connection is required for activation during installation.

Software Requirements

Software Requirements for the Dragon Client

The Dragon Client installer checks your system for minimum requirements. If the minimum requirements are not met, the installer will not install the client.

Operating system:

- Microsoft® Windows® 10 (including Professional and Enterprise), 32 bit and 64 bit
- Microsoft® Windows® 8.1, 32 bit and 64 bit
- Microsoft® Windows® 8 (including Professional and Enterprise), 32 bit and 64 bit
- Microsoft® Windows® 7, 32-bit and 64-bit
- Microsoft® Windows® Vista® Service Pack 2, 32 bit and 64 bit
- Microsoft® Windows XP® Service Pack 3, 32-bit only
- Microsoft® Windows Server 2008, Service Pack 1, Service Pack 2
- Microsoft® Windows Server 2008 R2, 32-bit and 64-bit
- Microsoft® Windows Server 2008 R2 64 bit Service Pack 2
- Microsoft® Windows Server 2012
- Microsoft® Windows Server 2012 R2

Software Requirements for the Nuance Management Server (NMS) Server

Operating system:

Choose one of the following:

- Microsoft® Windows Server 2008, Service Pack 1, Service Pack 2
- Microsoft® Windows Server 2008 R2, 32-bit and 64-bit
- Microsoft® Windows Server 2008 R2 64 bit Service Pack 2
- Microsoft® Windows Server 2012
- Microsoft® Windows Server 2012 R2 (64 bit)

Whichever operating system you choose, you should have all service packs installed, up to the most current one.

Microsoft PowerShell: Install the PowerShell feature along with Windows Server.

Microsoft .NET:

Microsoft .NET Framework 4.5

Web Server

- Internet Information Services (IIS). Version 7.0, 7.5, 8.0, and 8.5.

Internet Browser: Microsoft Internet Explorer 8, 9, 10, and 11 (free download at www.microsoft.com)

SQL Server

- SQL Server 2008
- SQL Server 2012
- SQL Server 2014e

Software Requirements for the Profile Optimizer Nodes

Operating system:

Choose one of the following:

- Microsoft® Windows 7
- Microsoft® Windows Server 2008, Service Pack 1, Service Pack 2
- Microsoft® Windows Server 2008 R2, 32-bit and 64-bit
- Microsoft® Windows Server 2012

Whichever operating system you choose, you should have all service packs installed, up to the most current one.

Microsoft .NET:

- Microsoft .NET Framework 4.5

SQL Server

Choose one of the following:

- SQL Server 2008
- SQL Server 2012
- SQL Server 2014

Windows Installer 3.1 If you are installing a *Speech Node* on a machine with any supported operating system other than Windows Server 2008, the installer requires that the machine have Windows Installer 3.1 or greater on it. (Windows Server 2008 automatically installs a later version of Windows Installer.)

Dragon Medical SDK Client You must install the latest version of the Dragon Medical SDK Client software on workstations where you plan to install *Speech Nodes*. This software is included on the *Nuance Management Server Software and Documentation DVD*.

Supported Virtual Desktop Infrastructure Applications and Terminal Emulators

- XenDesktop 7.1 and XenDesktop 7.6
- Wyse terminals

Supported non-EHR Applications

Once you have installed Dragon, you can use it to control the following applications using your voice:

- WordPad
- NotePad
- Microsoft® Word 2007, 2010 (32 & 64 bit), 2013 (32 & 64 bit), 2016 (32 & 64 bit)
- Microsoft® Outlook® 2007, 2010, 2013, 2016
- Microsoft® Excel® 2007, 2010, 2013, 2016
- WordPerfect® x5, x6
- Apache OpenOffice Writer 3.4
- Open Office Writer v3.1, 3.2
- Internet Explorer 8, 9, 10, 11 (11 supported when Enhanced Protective Mode is disabled)
- Rich Internet Application IE8, IE9, IE10, IE11
- Mozilla® Firefox® 8+
- Rich Internet Application Firefox® 12+
- Rich Internet Application Google Chrome 16+
- Windows Live Mail v15 & v16
- Mozilla® Thunderbird™ x3 and up
- Lotus Notes 8.5

Chapter 6: Network bandwidth requirements on DM Network Edition networks

This chapter contains information about configurations and recommendations for network usage and requirements. The information is based on several test environments set up by Nuance.

Network Bandwidth Requirements67

Network Bandwidth Requirements

Dragon clients are going to operate more fluidly in a faster network. And the more users you have dictating, the faster that network ought to be. The question every network planner needs to answer, however, is: How fast does the network *have* to be to ensure reasonable response time when users log in every morning and logins have to be authenticated on the *NMS Server*? What about when profiles have to be downloaded from the *Master User Profiles Server* for the first time? How can a network planner be sure the response time a user is going to experience remains reasonable?

You want to avoid setting up an underpowered network—where user output produces more network demand than the infrastructure allows. During initial deployment analysis, your Nuance representative should be able to help you take the right precautions up front to prevent any network issues. If network volume increases beyond the volume anticipated, or other traffic is generated on the network by additional applications being installed, you may need to modify your network to accommodate more volume.

Recommendations for Network Bandwidth

- Speed of the networks used to generate the response times shown in this document is 100 Mbps.
- Your network responses may differ as you scale the network and add applications. You need to know, especially, how much delay your EHR and other network applications such as email might produce.
- You should have a minimum 10 Mbps network if you are planning to have 100 *Dragon Clients* communicate with the *NMS Server*.
- To achieve the best response times, you should have a minimum 100 Mbps network if you have more than 100 *Dragon Clients*.

Considerations for Low Network Bandwidth Configurations

If you have configured your network to have low network bandwidth, using wireless and/or the public Internet, you can have the user select the **If your connectivity is slow, select this option to connect to the server without synchronizing your master profile** check box in the **Dragon Login** dialog box.

The *NMS Server* will still validate the user's license, but not communicate with the *Master User Profile Server* while the user is dictating. The user would work only with the local cache copy of the

profile. Making this choice restricts the number of instances of slow response to the processes of initially creating the profile and downloading the profile for the first time on a particular workstation.

To use this setting, the user must have already created a user profile and must have logged in to the server to access it at least once from the workstation being dictated on, so that a local cache copy of the profile exists on the *Dragon Client* workstation.

Note: This option is different from running in disconnected mode. In disconnected mode, users can see similar performance improvements because they are not interacting with the network. In order for a user to work in Disconnected mode, an administrator must check the **Allow disconnected mode** check box in the **Site** area in the Nuance Management Center. A user that is in Disconnected mode can use the Dragon client for 90 days. At that point, they must log in over the network.

Considerations for Network Bandwidth in Citrix Environments with vSync

The CPU load that vSync adds is negligible. vSync uses approximately 11 MB of memory per user dictation session.

Estimating Load and Traffic Generated in Citrix Environment Using vSync Feature

Another factor that might influence the function of your DM Network Edition network is setting the network up with a Citrix Server and using the vSync feature of DM Network Edition. vSync introduces a small delay that you can add to the delays stated earlier.

Load Created by 40 users Logging in on DM Network Edition Network with Dragon Using the vSync Feature

Load	Server idle	One user dictating	40 users with EHR and vSync	40 users with EHR typing	vSync load
CPU (%)	0.18 %	0.62 %	24.97%	22.21 %	2.76 %
Memory (MB)	718.4 MB	31.74 MB	1984.52 MB	1582.4 MB	402.09 MB

The table above shows that 40 users utilize 2.7% of the server’s CPU and that each user load is approximately 10 MB of memory (see the **vSync Load** column to the far right).

Other load information on *Dragon Clients* with vSync:

- With vSync loaded, but *Dragon Client* not active, no significant/measurable increase in CPU usage occurs.

- With vSync loaded, but *Dragon Client* not active, an increase in memory usage of approximately 4 MB per dictating user occurs.
- With vSync loaded and *Dragon Client* active, an additional increase in memory usage of approximately four MB per dictating user occurs (for a total increase of approximately 6-7 MB per user resulting from both vSync and dictating in Dragon).
- Sites using vSync with *Dragon Clients* would see about 130 KB of traffic for 3 minutes of dictation. This amount of traffic is similar to traffic generated by instant messaging.

Estimating Dragon Response Times When user Actions Access NMS Server

Since the response times vary depending on network traffic at any given moment, it's a good idea to take a look at the response times that occurred in Nuance tests on *Dragon Clients* in *DM Network Edition* network.

Measurements of Dragon Client Response Times on Network

Response Time of user Logging In While Increasing Numbers of Other users Log In, Retrieve Dragon Settings, Change Dragon Settings, or Log Out of the Dragon Client on 100 Mbps Network

Total Active users	Minimum Login time	Maximum Login time	Mean Login time	Range	Mode	Median
No users taking other actions	2 seconds	3 seconds	2.33 seconds	1 second	2 seconds	2 seconds
100 to 150	2 seconds	3 seconds	2.66 seconds	1 second	3 seconds	3 seconds
200 to 300	3 seconds	4 seconds	3.33 seconds	1 second	3 seconds	3 seconds
400 to 600	3 seconds	4 seconds	3.77 seconds	1 second	4 seconds	4 seconds
600 to 900	4 seconds	5 seconds	4.22 seconds	1 second	4 seconds	4 seconds
800 to 1,200	4 seconds	5 seconds	4.55 seconds	1 second	5 seconds	5 seconds
1,000 to 1,500	4 seconds	5 seconds	4.77 seconds	1 second	5 seconds	5 seconds

Bear in mind that these numbers reflect the response time on a 100 Mbps network. Your network might need to be at least as fast if you are planning on having this many or more users simultaneously taking actions that require accessing the *NMS Server* over the network.

Although you can configure up to 1,000 users on each *NMS Server*, typically fewer than 25% of those professionals would be actually dictating at one time and only a percentage of those users (perhaps 3% to 5% of your total user population) would be logging in or otherwise interacting with the server at the *exact same time*.

You should also take into consideration that any other applications, including your EHR application, generate their own share of traffic.

Approximate Sizes of Data Transferred Over DM Network Edition Network and Effect on Network Load

Another tool you can use to determine the amount of traffic and consequent delays you might experience on your *DM Network Edition* network is the amount of data transferred over the *DM Network Edition* network when the user takes particular actions on the client. Refer to the table that follows.

Sizes of Data Transferred over the Network for Specific Client Actions and Approximate Response Time per user on 100 Mbps Network

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
User creates new user profile	Profile created on client and stored on <i>Master User Profile Server</i> . Occurs only once for any particular user.	300 MB +86 MB for additional vocabulary + 60 MB for additional dictation source	Dragon client	Master User Profile Server	Access network at user open or close only
		About 10 MB if 2000 words and 200 auto-texts present	NMS Server	Dragon client	
Dragon saves profile after dictation	Small profile changes copied from client to <i>Master User Profile Server</i> .	About 105 MB of raw data after 10 minutes of dictation	Dragon client	Master User Profile Server	Saving user also saves local acoustics
User logs in to the NMS server through the client for the first time	User profile copied from <i>Master User Profile Server</i> to client; settings sent from <i>NMS Server</i> to client.	175 MB	Master User Profile Server	Dragon client	
		About 2 Mb for a user without words, commands, or auto-texts	NMS Server	Dragon client	
User logs in to the client after the first time	(Client is accessing user profile from cache) <i>NMS</i>	2 MB	NMS Server	Dragon client	

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
(after initial log in, but not first time after optimization)	Server passes settings to client.				
User dictates in Disconnected mode, then reconnects to network	Data from Disconnected mode sessions and changed settings copied from local cache to <i>Master User Profile Server</i> .	Less than 1 MB + size of the dictation audio files (10 MB/min)	Dragon client	Master User Profile Server	Allow Disconnected Mode
User logs in for the first time after an optimization completes	An updated user profile copied from <i>Master User Profile Server</i> to client.	240 MB	Master User Profile Server	Dragon client	
user dictates and makes corrections	Incremental changes to user profile copied to <i>Master User Profile Server</i> .	10 MB per minute of audio	Dragon client	Master User Profile Server	Conserve server archive size
user logs out of Dragon	User profile changes copied to <i>Master User Profile Server</i>	1 MB + 10 MB/min of DRA (audio) files	Dragon client	Master User Profile Server	
user carries out Training	User profile changes copied to <i>Master User Profile Server</i> .	15 MB	Dragon client	Master User Profile Server	
	After ACO completes, a Speech Node creates Dragon Log file and Speech Node Service log file.	6 MB to 10 MB (3 to 5 MB per log)	Speech Node	NMS Server	Copy Dragon log to the network
	After ACO completes, a Speech Node saves the optimized profile to the Profile Server	650 MB	Speech Node	Master User Profile Server	
	NMS Server creates foundation service log file.	10 MB	NMS Server	NMS Server Database	

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
	When LMO completes, the Speech Node transfers the optimized LM to the Master User Profile Server.	15 MB	Speech Node	Master User Profile Server	
	After medications are integrated into each user's LM, the Speech Node transfers the LM to the Master User Profile Server.	15 MB	Speech Node	Master User Profile Server	

Process Producing Greatest Network Load

Network load is affected most when users are in the process of creating a new user profile. During this process, *DM Network Edition* transfers 70 MB of data from the client to the *Master User Profile Server*. *DM Network Edition* transfers additional data for additional vocabularies and dictation sources. Usually, a user has only one vocabulary, such as **Cardiology**, based on a specialty. At most, a user might have two vocabularies.

Also, *DM Network Edition* transfers additional data if the user has more than one dictation source/device, whether that means multiples types of microphones, multiple types of recorders, or both. users likely have a maximum of two types of microphones and/or two types of recorders.

Since each user creates a new profile only once, this activity does not have any on-going effects on the function of your network.

Traffic Produced by Downloading a user's User Profile

For each workstation where a user dictates, *DM Network Edition* must download the user's user profile to that machine *only the first time* the user uses *Dragon* on that workstation. That initial download is 70 MB of data; however, that download occurs only once for a user per workstation that the user dictates on. This action occurs only on first using the workstation and, later, after an optimization process completes. The frequency of the ACO processes can be scheduled to reduce their impact.

Scheduling ACO Processes and Managing Master User Profile Downloads to Clients

Usually user profile data on the *Dragon Client* workstation is newer than the data stored in the user profile on the *Master User Profiles Server*.

The one time that the data on the server is newer is immediately after an optimization process occurs.

After ACO occurs, the client must download a copy of the master user profile for the user during the next login to the *NMS Server*. Since the administrator can schedule how often an ACO occurs for each individual user (weekly, monthly, bi-monthly, and so on), that administrator can predict how often profiles will be affected and reduce how frequently users experience delay during the log in process.

Effects of LMO Processes and Medication/Command Updates

While ACO processes move significant data, LMO and Medication or Command Update processes move very small amounts of data and need to move only one copy of the data, so those numbers are not per user profile, but per update, and have no significant effect on the network.

Traffic Produced by On-Going Daily Actions/Dictation

Once the user profile has been established for a user, a local cache copy of the profile is on the *Dragon Client* workstation. Then, on a typical day, when a user opens his or her user profile on the *Dragon Client* in a *DM Network Edition* network, the client opens that user profile from the local cache, rather than over the network.

Only logging in, logging out, and modifying/retrieving *Dragon* settings forces the client to go out over the network to interact with the *NMS Server*. The size of the settings being transferred is equivalent to the size of a small email. How fast a profile opens is mostly based on the speed of the CPU on the machine where the *Dragon* client is installed.

The amount of network traffic that flows from the workstation to the *Master User Profile Server* on a typical day depends mostly on how much the user dictates and corrects recognized text. Nuance quantifies the data transferred during this process using this rule of thumb:

Every 1 minute of dictation transfers approximately 10 MB of data.

Large Distributed Networks Can Deploy Multiple Distributed Master User Profile Servers

In a large distributed network configuration, you can have multiple separate distributed *Master User Profile Servers* so that clients can have quick local access to the user profiles. In a situation like this, a connection between the client machine and the local *Master User Profile Server* can be much slower

than the network connection between distributed sites; for instance, 10 Mbps might be an adequate speed for the local area network, whereas 100 Mbps might be required for the wide area network to connect to other sites.

Large Distributed Networks Can Deploy a Load Balancing Switch to Manage Multiple NMS Servers

In a large distributed network configuration with the NMS Server on your network, you can deploy a network traffic switch to manage load balancing. The network traffic switch (available through F5 and similar manufacturers) can send a request to each server to query its status, tag a server as down if it does not respond, and reroute the traffic to other servers.

- Logging in
- Retrieving Dragon settings
- Modifying Dragon settings
- Logging out

Test Environment for Measuring Estimating Load and Traffic Generated in Citrix Environment Using vSync Feature

Hardware Item	Citrix Server Hardware and Software	Dragon Client Hardware and Software
Memory	3827 MB	1.46 GB
Processors	4 Intel(R) Xeon (R) CPU E5530 @ 2.40 GHz @ 2531 MHz	Intel Xeon 2.40 GHz
Client Software/Version	<i>Dragon Client vSync</i> Version 10.1	<i>Dragon Client</i> Version 10.1
Test Software	EdgeSight Version 5.2.3012.0	SilkTest Version 7.6.0., v5.02
Server Software	Citrix Presentation Server Version 4.6.3600.0 (SP 2006.10), Platinum Edition	Citrix Program Neighborhood Version 11.0.5357 (256-bit SSL)

Hardware Item	Citrix Server Hardware and Software	Dragon Client Hardware and Software
Operating System	Windows Server 2003 R2 Standard Edition SP2 32-bit	
Other Information		40 Virtual machines running <i>Dragon Client</i> software.

Index

A		E	
accent in speech		errors in recognition	
effect on how often to run ACO	18	effect on how often to run ACO	18
ACO process	17		
defined	17		
how it differs from LMO	17		
how often to run	18		
maintenance	19		
minimum recommended frequency	18		
recommended frequency	18		
adaptation			
defined	17		
B		H	
basic components of the Dragon solution	7	hardware/software requirements	
Bluetooth wireless microphone support	63	master user profiles storage	44
		hardware/software requirements summary	
		network configurations	46
C		I	
calculation		improved recognition	
number of minutes for ACO	22	how ACO results affect	21
number of nodes required	21	IP switch	
percentage of providers dictating	23	load balancing on NMS Servers	27
space for local cache user profiles	15		
space for master user profiles	13		
calculation example			
space for master user profiles	14		
D		K	
databases		knowledge	
storage space for NMS Server	25	required to plan network	3
dictation volume			
effect on how often to run ACO	18		
effect on how often to run LMO	20		
how often to run LMO based on	20		
Dragon			
network components	7		
overview	7		
Dragon logs			
how often purged	26		
duplex settings on NICs	27		
		L	
		large configuration	
		hardware and software details	42
		link speed settings on NICs	27
		LMO process	17
		defined	17
		how it differs from ACO	17
		how often to run	20
		load balancing	
		multiple NMS Servers	27
		local cache user profiles	
		how ACO results integrate with	21-22
		space requirements calculation	15
		log files	
		storage space requirements	26
		logs	
		how often purged	26
		M	
		master user profiles	
		space requirements calculation	10

master user profiles machine
hardware/software requirements 44

medium configuration
hardware and software details 40

N

network configurations
hw/sw requirements summary 46
summary table of servers 49

network equipment speed 27

network interface cards supported 27

network planning
background required 3
knowledge required 3
skills required 3

network switch settings 27

network traffic switch
load balancing 27

NMS Servers
balancing load on multiple 27

O

optimizations
number of nodes required 21

P

pronunciation
processes for learning 17

providers
how often to run LMO for experienced 18
how often to run LMO for new 20
Speech Nodes required by number of 24

R

recognition improvement
how ACO results integrated for 21

S

server installations
storage space requirements 26

settings
network interface cards 27
network switches 27

skills
required to plan network 3

small configuration
hardware and software details 39

Speech Node service logs
how often purged 26

Speech Nodes
calculating number required 21
number based on total providers 24
role in ACO and LMO processes 17

speech patterns
processes for learning 17

speed of network equipment 27

storage space
requirements for log files 26
requirements for NMS database 25

Storage space
Server installations 26

storage space requirements
calculation for local cache user profiles 15
calculation for master user profiles 10

summary table
Network configurations 49

System requirements
Dragon 10 53

T

tuning acoustic and language models 17

U

user profiles
number of nodes for optimization 21
optimization 17
space requirements calculation 10

V

vocabularies
effect on how often to run LMO 20

W

words
effect on how often to run LMO 20