Dragon Medical 360 | Network Edition

2.1

Planning and Deployment Guide

Copyright

2013. Dragon Medical 360 | Network Edition, 2.1.

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 © 2002-2013 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 noninfringement.

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 360 | Network Edition software contains configurable password features that, when used properly, provide a high degree of protection.

L-3690

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 360 | Network Edition 1

Introducing Dragon Medical 360 | Network Edition 2

Planning an installation of Dragon Medical 360 | Network Edition 3

What's new in DM360 Network Edition 4

What's new in DM360 Network Edition, 2.1 4

Chapter 2: Dragon Medical 360 | Network Edition network requirements 5

Dragon Medical 360 | Network Edition Overview 7

Dragon Medical 360 | Network Edition components 8

Nuance Management Server components 8

Nuance Management Server in the Cloud 9

Profile Optimizer components 10

Dragon Medical Clients 10

Master user profiles directory 10

Nuance hosted servers 11

Call Forwarder 11

Who works on the network 11

Determining space requirements for user profiles 12

Understanding master user profiles and local cache user profiles 12

Determining storage space required for master user profiles 13

Determining storage space required for local cache user profiles on client machines 16

Deciding how often to update acoustic and language models 18

Understanding ACO and LMO processes 18

Defining ACO/LMO processes and their impact on recognition 18

Factors in how often to run ACO processes 19

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 providers 23

Estimating storage requirements for Nuance Management and Profile Optimizer Server databases 25

About Nuance Management Server (NMS) and NMS SQL database ... 25
About Profile Optimizer Server and Profile Optimizer 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 ... 27
Ports to open for clients, servers, and hardware firewalls ... 28
Summary: Ports to open on Clients, Servers, and hardware firewalls ... 30
Chapter 3: Configuring DM360 Network Edition networks that use physical servers ... 33
Using physical servers in a DM360 Network Edition network ... 34
Configuring a DM360 Network Edition network for up to 100 providers ... 34
Configuring a DM360 Network Edition network for 100 to 1000 providers ... 34
Configuring a DM360 Network Edition network for more than 1000 providers ... 35
System requirements for small networks with physical servers ... 40
System requirements for medium networks with physical servers ... 41
System requirements for large networks with physical servers ... 43
System requirements for master user profile computers ... 44
Hardware and software requirements ... 45
Optional hardware ... 45
Optional software ... 46
Summary: System requirements for DM360 Network Edition networks that use physical servers ... 47
Summary: Configuration options for DM360 Network Edition networks that use physical servers ... 50
Chapter 4: Configuring DM360 Network Edition networks that use virtual servers ... 52
Using virtual servers in a DM360 Network Edition network ... 53
Using Virtual Machines for the NMS Server or NMS Call Forwarder, the Profile Optimizer, and databases ... 53
Using Virtual Machines for Profile Optimizer Speech Nodes ... 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 ... 57
Chapter 5: System requirements for the Dragon client and user profiles on a Profile Optimizer Server ... 59
System requirements for the Dragon Medical Client ... 60

System requirements for Master User Profiles on a Profile Optimizer Server ... 62

Chapter 6: Network bandwidth requirements on DM360 Network Edition networks ... 63

Network Bandwidth Requirements ... 64

Recommendations for Network Bandwidth ... 64

Considerations for Low Network Bandwidth Configurations ... 64

Considerations for Network Bandwidth in Citrix Environments with vSync ... 65

Estimating Dragon Response Times When Provider Actions Access NMS Server ... 67

Approximate Sizes of Data Transferred Over DM360 Network Edition Network and Effect on Network Load ... 68

Summary of Test Environments ... 73

Index ... 75

Chapter 1: Preparing to install Dragon Medical 360 | Network Edition

Introducing Dragon Medical 360 | Network Edition 2
Planning an installation of Dragon Medical 360 | Network Edition 3
What's new in DM360 Network Edition 4

Introducing Dragon Medical 360 | Network Edition

Dragon Medical 360 | Network Edition is a revolutionary product for Healthcare Enterprise environments; providing a central server (the *Nuance Management Server*) to manage multiple local or remote speech recognition client machines where healthcare providers dictate. DM360 Network Edition takes advantage of distributed processing across servers and workstations for greater efficiency.

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

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

- Plan a deployment of a *Dragon Medical 360 | Network Edition Network*.
- Determine the number of servers and other equipment you will need.
- Prepare equipment for the DM360 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 Medical Clients* that work with both the *NMS Server* and the other servers and components of the network.

Notes:

- Some DM360 Network Edition features might have changed since this manual was printed. A current version of this book is always available on Nuance's documentation portal, in PDF format. Contact your Nuance representative about accessing that portal to retrieve the latest copy.
- The NMS license includes one Per Processor License for the SQL 2008 OEM Standard Edition. When you purchase an NMS license, the cost of the license is for the NMS software and not the SQL license. Nuance provides the SQL license to you at no extra cost.

Planning an installation of Dragon Medical 360 | 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. To plan the network and install the *Dragon Medical 360 | Network Edition* server and client software, you should have adequate skills and experience to:

- 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

Note:

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 360 | 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 DM360 Network Edition, see the DM360 Network Edition *Installation Guide* or the *Nuance Management Server Administrator Guide*

What's new in DM360 Network Edition

What's new in DM360 Network Edition, 2.1

DM360 Network Edition, 2.1 includes the features, enhancements, and bug fixes from all previous service packs and hot fixes, plus new features.

For details about what's new in DM360 Network Edition, 2.1 for administrators, see the following topics in the Nuance Management Server Administrator guide:

- What's new in Nuance Management Server - All products
- What's new in Nuance Management Server - DM360 Network Edition

Data objects are stored in (Nuance Management Server) NMS in the cloud (in the Nuance data center) instead of an on-premise NMS. These objects include:

- User accounts
- Text and Graphics/auto-texts
- Custom words
- Custom command sets

To learn more about NMS in the cloud, see the 'Introducing NMS in the cloud' section in the Nuance Management Server Administrator guide or the DM360 Network Edition, 2.1 release notes.

Chapter 2: Dragon Medical 360 | Network Edition network requirements

You can configure your Dragon Medical 360 | Network Edition network using the *Nuance Management Server* in several ways.

If you have a small enterprise, you can install all server components of the system on a single server machine. DM360 Network Edition servers are flexible enough that if you have a larger organization, you can scale the system by installing server components on separate machines to balance the workload.

In addition, you can deploy multiple *Profile Optimizer Speech Nodes* to handle a high volume of user profile optimization for a large and prolific staff.

If you are using the Nuance Management Server in the cloud (Nuance data center) to store data objects, you need only to deploy some of the server components (the Profile Optimizer, Speech Nodes, and possibly the Call Forwarder) within your organization's network.

This chapter discusses all of the DM360 Network Edition components and helps you determine the exact needs for your organization's network including the space you need for master user profile storage, the number of *Speech Nodes* you need, the space required for the *NMS Server* and *Profile Optimizer Server* databases, and the type of network equipment and switch settings recommended.

Dragon Medical 360 | Network Edition Overview 7

Dragon Medical 360 | Network Edition components 8

Determining space requirements for user profiles 12

Deciding how often to update acoustic and language models 18

Determining number of Speech Nodes required for optimizations 20

Estimating storage requirements for Nuance Management and Profile Optimizer Server databases 25

Recommended network and switch settings 27

Using Network traffic switch for load balancing 27

Storage hardware requirements for master user profiles 27

Ports to open for clients, servers, and hardware firewalls 28

Dragon Medical 360 | Network Edition Overview

Nuance Management Console Help for Dragon Medical 360 | Network Edition and the *Nuance Management Server (NMS Server)* help you meet your expanding medical organization's enterprise level needs. You use the *NMS Server* to manage multiple *Dragon Medical Clients* from a single central server.

Network setup is now straightforward and efficient when you use the *NMS Server* to apply consistent settings to many workstations running *Dragon Medical Client* throughout the medical facility, changing multiple settings on the *NMS Server* with the click of a mouse on its *Nuance Management Console*.

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

- Create multiple sites where *Dragon Medical 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 healthcare provider (**Physician Author** and **Non-Physician Author**) 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 Medical Client* workstation.
- Initiate, schedule, and monitor acoustic/language model adaptation 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.
- Set a common Data Distribution Directory for all machines at one or more groups.
- Create user accounts for healthcare providers, 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 Medical Clients* .

Dragon Medical 360 | Network Edition components

Dragon Medical 360 | 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
- *Profile Optimizer* components
- *Dragon Medical Clients*
- Master user profiles directory (on any server of the network or on its own machine)
- Nuance hosted servers or the Call Forwarder

Nuance Management Server components

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

- **Nuance Management Server (NMS Server)** — Main Web service that drives DM 360 | Network Edition and helps it interact with external clients and servers by carrying out several vital tasks involving allowing user access, determining settings on servers and clients, and storing information about user preferences and transcription formatting. In addition, the *NMS Server* works with the *Profile Optimizer Server* to ensure utilization of information about a particular speaker's voice and language models. 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 Console** — An interface to not only the *Nuance Management Server* for managing set up of all *Dragon Medical Clients* on the network, but to many functions of the *Profile Optimizer Server* as well.
- **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 Console* on the same machine; however, you have the option of installing the database on the same or a separate physical server on your network.

The following conceptual illustration shows the components in the Dragon Medical 360 | Network Edition.

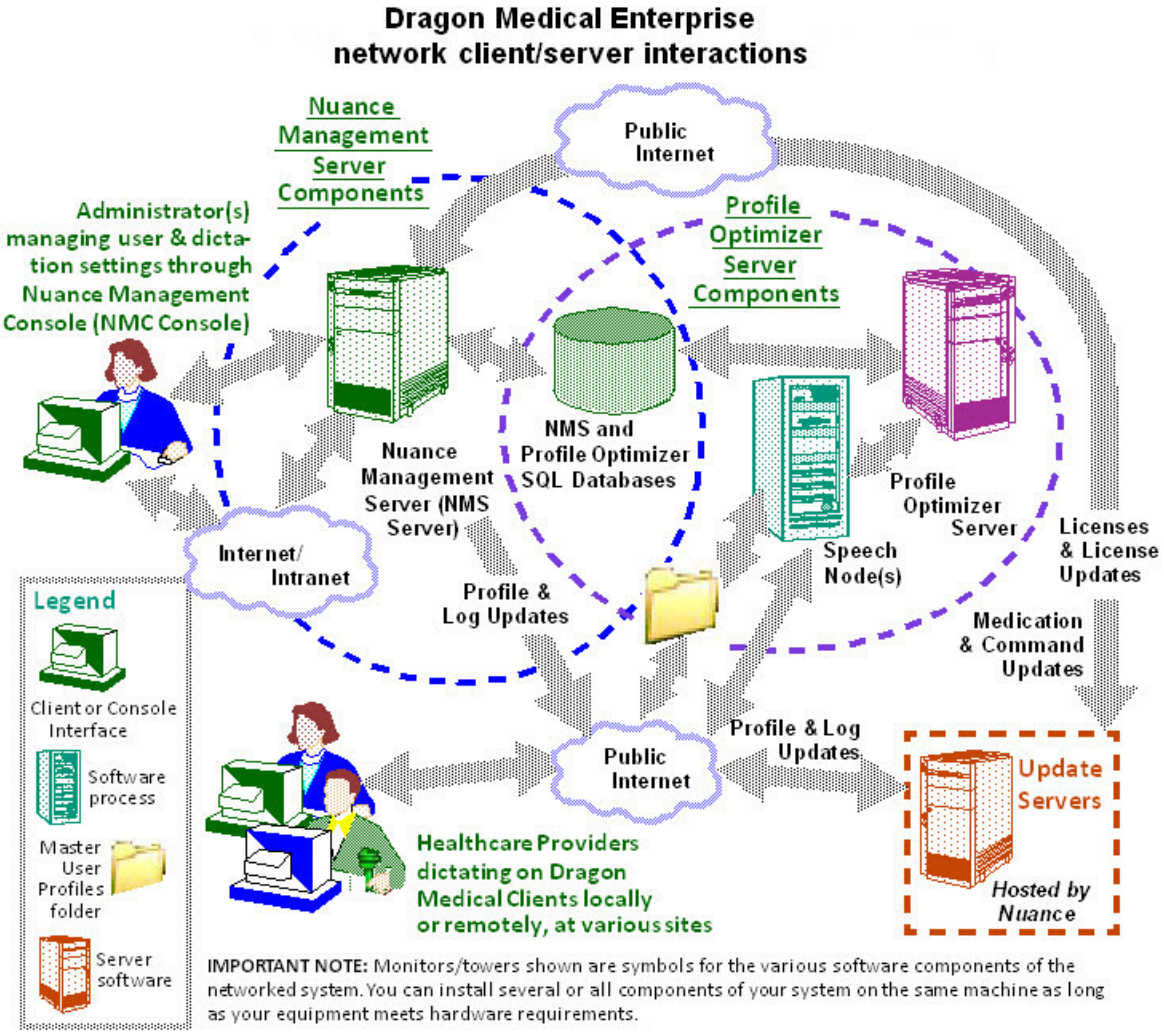

Nuance Management Server in the Cloud

If your end users need to share profile customizations with Nuance cloud-based applications, you do not install the Nuance Management Server (NMS) components locally; instead you use the existing NMS in the Nuance data center. This configuration is referred to as "NMS in the cloud." Nuance cloud-based applications include Dragon Medical 360 | Direct and third-party applications that use Dragon Medical 360 | SpeechAnywhere Services.

The profile customizations that can be shared using NMS in the cloud are:

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

For example, if an administrator creates an auto-text in the Nuance Management Console, users can use that auto-text in Dragon Medical 360 | Network Edition, Dragon Medical 360 | Direct, and applications that use Dragon Medical 360 | SpeechAnywhere Services.

Using NMS in the cloud is optional, and is required only if you plan to share customizations with Nuance cloud-based applications. For more information on Nuance cloud-based applications, consult your Nuance account executive.

If you use NMS in the cloud, you still need to install the Profile Optimizer components in your network. You also need to install the NMS Call Forwarder if you want to use LDAP or Active Directory to authenticate Dragon Medical 360 | Network Edition users.

Profile Optimizer components

Profile Optimizer components of the network include:

- **Profile Optimizer Server** — Server that manages the iterative learning and integration of each user's unique speaking/pronunciation patterns by scheduling two processes, Acoustic Model Optimization (often referred to as simply ***Acoustic Optimization*** or ACO) and Language Model Optimization (LMO). The ACO process integrates information about how the individual dictating speaks and pronounces words while the LMO process integrates 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.
- **Profile Optimizer Speech Nodes** — Modules that carry out the Acoustic Optimization (ACO) and Language Model Optimization (LMO) for each dictating healthcare provider.
- **Profile Optimizer Database** — Database that stores the schedules for the Acoustic Optimization (ACO) and Language Model Optimization (LMO) tasks. If you are running the NMS in your network (as opposed to using NMS in the cloud), you usually place this database on the same machine where you created the *NMS SQL Database*.

Dragon Medical Clients

Dragon Medical Client — Receives and interprets healthcare provider dictation, producing text corresponding to each physician's speech. The *Dragon Medical 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 providers. A single central storage location applies to a particular site in DM 360 | Network Edition. A master user profile contains the accumulated speech information for a particular healthcare provider from all the *Dragon Medical* client workstations where that person has dictated, 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, over time, the *Profile Optimizer Server* and its *Speech Nodes* learn the particular user's speech patterns and integrate them into that individual user profile. These user profiles are called *roaming user profiles* in other editions of *Dragon Medical*, 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 Microsoft Windows user profiles.)

Nuance hosted servers

Nuance hosts servers and services at its headquarters that your DM 360 | Network Edition network interacts with:

- **Update Server** — Hosted by Nuance on the Internet. Automatically downloads various updates to the *NMS Server* over the Internet.

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 360 | Network Edition network.

Call Forwarder

If you will be using LDAP or Active Directory to authenticate users and if you are using NMS in the cloud, you need to install a small server application called the NMS Call Forwarder. The call forwarder tests user credentials and sends authentication requests to the cloud NMS to retrieve data.

Who works on the network

Two types of people work on the network:

- **System Administrators** — Manage the entire system through the *Nuance Management Console*.
- **Healthcare Providers** — Dictate using *Dragon Medical Clients*.

Determining space requirements for user profiles

User profiles reside within your network, whether you are running NMS inside your network or in the cloud.

To determine how much storage space your network requires to manage user profiles for all of your dictating healthcare providers, 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 Medical* uses to help recognize the individual provider'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 provider logs into *Dragon Medical* on a client machine, *Dragon* makes a copy of that provider'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 provider logs out of the *Dragon Medical 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 provider logs into the *Dragon Medical 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 Medical network requires to store *master* user profiles. How much storage space is required depends on how many providers you have and other specific data about those providers. 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 Medical Client* machine requires to store *local cache* user profiles for all providers 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 healthcare providers are dictating with *Dragon Medical* throughout the entire organization?
- How many providers have an extra vocabulary? You should add an extra vocabulary for each provider who practices in more than one medical specialty, as well as for any provider 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 providers 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 healthcare providers do you anticipate adding to the population of providers dictating with *Dragon Medical* in the next year?

Notes about acoustic models and user profiles

- BestMatch IV and BestMatch V acoustic models on a dual-core machine are about 36 MB larger than BestMatch III acoustic models (30 MB).
- When you upgrade a profile to DM360 Network Edition 2.1, the Acoustic optimizer processes several .dra files. The processed .dra files are larger than the original .dra files.
- When you upgrade a profile to DM360 Network Edition 2.1, 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 Healthcare Providers Dictating					
Total Number of Providers (N)	**Providers with one extra vocabulary (V)**	**Number of extra audio input device(s)(A)**	**Number of Providers Added Annually (P) (optional)**	**Acoustic Model Optimization Data Storage (M)**	**Multiply Total by 2.05 to Allow for Backups**
Allow 60 MB for each provider.	Allow 20 MB for each extra vocabulary (Add another vocabulary when user also dictates outside Medical system into email or office apps).	Allow 60 MB for each extra audio input device (dictation source).	Add 53 MB for each provider, allowing each an extra audio input device and an extra vocabulary	Allow 1 GB of space for each provider times the number of audio input devices (dictation sources) the provider 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 x 60 MB)	+ (V x 20 MB)	+ (A x 60 MB)	+ (P x 53 MB)	((N + A + P)(1 GB))	x 2.05
	+	+	+	+	x 2.05
TOTAL:					

Example: Calculating storage space required for master user profiles

Number of Healthcare Providers Dictating					
Total Number of Providers (N)	**Providers 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 providers	12 with one extra vocabulary and 3 with two extra vocabularies	Each provider has a headset and a PowerMic II. First microphone already included in original space; second considered additional.	Planning to add 5 new providers in the coming year.	Multiply number of providers 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.

Number of Healthcare Providers Dictating					
Total Number of Providers (N)	**Providers 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**
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)	(200 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* and *Profile Optimizer* software, 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 Medical* configurations that have more than 10 users.

Note:

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 Medical Client* workstation to make a local copy of the master user profile for the client to use when the provider dictates.

Notes about acoustic models and user profiles and Dragon Medical 360 | Network Edition 2.1

- BestMatch IV and BestMatch V acoustic models on a dual-core machine are about 36 MB larger than BestMatch III acoustic models (30 MB).
- When you upgrade a profile to Dragon Medical 360 | Network Edition 2.1, 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 360 | Network Edition 2.1, 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 Healthcare Providers Dictating				
Number of Providers Dictating on Client (C)	**Number of Providers 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 provider.	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 provider times the number of audio input devices (dictation sources) the provider has.
((C x 60 MB)	+ (V x 20 MB)	+ (A x 60 MB)	+ (C + V) (10 MB)	+ (C + A)(500 MB))
TOTAL:				
Example: Allocating space for 10 dictating healthcare providers:				
10 providers	2 with one extra vocabulary and 1 with two extra vocabularies	All providers have a headset and a PowerMic II. The first microphone is already included in original space and the second is considered extra.	Every user has one vocabulary with a model to be optimized, so add 10 MB for every provider.	Multiply the number of providers by 2 because each has 2 microphones; then multiply that total by 240 MB.

Number of Healthcare Providers Dictating				
Number of Providers Dictating on Client (C)	**Number of Providers 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)**
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 Console*, you can set the amount of disk space to be reserved for archived corrections for user profiles on each workstation running a *Dragon Medical 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 healthcare provider. *Dragon Medical Enterprise* tunes each provider's master user profile with data from his or her unique acoustic model and language model.

Understanding ACO and LMO processes

The *Profile Optimizer* manages the aspect of speech recognition that involves learning each healthcare provider'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 *Profile Optimizer Server* updates that healthcare provider'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 healthcare provider. Research has shown that running both ACO and LMO processes regularly reduces the relative word error rate for the provider by 10 to 30% and can produce even greater improvements in recognition accuracy for providers 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 healthcare provider?

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 provider 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 Medical*, 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 provider's speech. The model incorporates how the provider 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 effect the speed of recognition during dictation on *Dragon Medical Clients*.

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

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

To get the most out of the ACO process, you should schedule the process to take place every time the provider has dictated 5 hours of audio. A medical practitioner 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 provider 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 provider is new to *Dragon Medical*. After the provider has been dictating for a month or has dictated 5 hours of total audio, the provider 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 provider dictates 5 hours of new audio.

If the *Profile Optimizer Speech Node* does not find enough speech data for that provider 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*, then the *Speech Node* needs to carry out a separate ACO process for each of those devices

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

Running a maintenance ACO process periodically keeps the provider'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 providers whose speech *Dragon Medical* has difficulty recognizing or providers who have ***ever*** had their profile becomes corrupted.
- At least every week (but preferably every day) for providers 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 providers 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 provider or group of providers generates
- How often you add new words to vocabularies
- How often you would like the *Profile Optimizer* 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 from the *Update Server*, you do not take any other action. The *Profile Optimizer 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 360 | Network Edition that updates and maintains acoustic and language models is the *Profile Optimizer Server*. DM360 Network Edition includes a *Profile Optimizer Server*. The Profile Optimizer Server always resides within your network, whether you are running NMS inside your network or in the cloud, 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 provider are called *acoustic model optimization* and *language model optimization*.

Before you try to determine how to include a *Profile Optimizer 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 healthcare providers, you can then estimate how many *Profile Optimizer Speech Nodes* your network requires to regularly integrate and optimize the information *Dragon* learns about each provider'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 provider'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 healthcare providers 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 *Update Server* 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 provider gives you a good starting point to determining how many *Profile Optimizer Speech Node*s 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 provider has logged in for the first time and has been speaking into the microphone, after the provider finishes dictating and logs out, the *Dragon Medical Client* sends data about the provider's speech to the corresponding master user profile on the network. If that provider has spoken for at least 15 minutes, the *Profile Optimizer Speech Node* then carries out an ACO process on the acoustic model for that provider and updates that master user profile accordingly. Later, when the provider logs on to dictate again, the *Dragon Medical Client* retrieves the optimized master user profile from the network and places it in its local cache. The provider should then see improved recognition.

After the provider again dictates, then logs out of the session, the client again sends data about the provider's speech to the master user profile on the network. The provider 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 provider even knowing about them.)

The first week that the provider 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 provider dictates, until the provider'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 provider 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 provider during that month, the node terminates the process and waits until the next ACO that is scheduled.

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

Number of minutes required for adaptation processes

Days of Provider Dictation	Stage of Provider "Training"	Minutes of Dictation per Provider Before ACO	Number of Providers One Speech Node Supports	Time Required to Tune Acoustic Model for Number of Providers with One Speech Node
First few days	First few days	15 minutes	_80 providers	80 providers X 15 minutes each = 1200 minutes in one day
End of 1st Week	Early dictation	50 minutes	_25 providers	25 providers X 50 minutes each = 1200 minutes per day
End of 2nd Week	More dictation	100 minutes	_12 providers	12 providers X 100 minutes each = 1200 minutes per day
End of Month		200 minutes	_10 providers	10 providers X 200 minutes each = 1200 minutes
Upgraded and in Maintenance mode	Upgrade training and ongoing dictation	5 hours (300 minutes)	125 providers x 30% dictating at one time = 40 providers	40 providers X 30 minutes per day = processes acoustic models at 1200 minutes per day; two nodes can share ACO load during upgrade process

If a healthcare provider uses more than one audio input device (dictation source), such as one USB microphone and one *PowerMic II*, 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 healthcare provider. *Almost* the same, but not exactly the same, because a provider who divides time between two devices probably spends less time dictating with each single device than another provider spends dictating with only one device.

A single *Speech Node* can handle ACO and LMO for only 40 to 80 providers if they are all just dictating for their first few weeks, but that same node can handle 125 providers who have been dictating long enough to require an ACO be scheduled only once a month.

Probably not all of your providers 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 providers a single node can process ACOs for at various levels of dictation. If you add up the number of providers 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 providers are going to be dictating at the same time. For instance, if you have 240 providers, 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 providers 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 providers dictate at any one time.

Number of Speech Nodes required in your installation

Time Provider Has Dictated	Approx. Number of Providers 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 One Month (Maintenance)		5 hrs	X 30	
TOTALS				_____ x _____% /1200 = _____ Nodes

Recommendations about number of providers

A good rule of thumb for how many providers a single *Speech Node* on a dual core machine can process ACOs for is 125 providers who are in maintenance mode per processor core beyond the first core. (The first core is reserved for all other server processes of the DM360 Network Edition network.) So, when the *Speech Node* is on the same machine with the *Profile Optimizer Server*, you can have up to 125 providers in maintenance mode as long as the machine is dual core.

If you have a dual core machine dedicated to *Speech Nodes*, you can create two virtual machines, one for each core, 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 Providers	Interval between Acoustic Model Optimizations (ACOs) in Months					
	1 mth recommended	2 mths	3 mths	4 mths	5 mths	6 mths minimum
0 – 120	1	1	1	1	1	1
121 – 240	2	1	1	1	1	1
241 – 360	3	2	1	1	1	1
361- 480	4	2	2	1	1	1
480 – 600	5	3	2	2	1	1
601 – 720	6	3	2	2	2	1
721 – 840	7	3	3	2	2	2
841 – 960	8	3	3	2	2	2
961 – 1080	9	4	3	3	2	2
1,081 – 1,200	10	4	4	3	2	2
1,201 – 1,320	11	5	4	3	3	2

If you have one or more *Speech Nodes* on the same machine with both the *Nuance Management Server* and the *Profile Optimizer Server*, note that the machine must have at least a quad core processor.

That quad core processor lets you can install up to three virtual machines for *Speech Nodes*on the physical server or workstation , one virtual machine per *Speech Node* for each additional core beyond the first core. The number of providers you can process with those three nodes is from 241 to 360 (up to 120 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 providers to between 480 and 600.

If you have the *Speech Nodes* on their own machine, separate from the servers, you can have four virtual machines on a quad core physical server, each running a single node, and all those *Speech Nodes* can together run ACOs every two months for from 1,081 to 1,200 healthcare providers.

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

Estimating storage requirements for Nuance Management and Profile Optimizer Server databases

To estimate how much disk space databases for the *Nuance Management Server* and *Profile Optimizer* require, you need to understand the kind of information each one contains.

About Nuance Management Server (NMS) and NMS SQL database

The *NMS SQL Database* stores information about the customer account for the organization as well as sites, users, groups, and licenses. If you are running the NMS within your network, the NMS SQL Database is within your network as well. If you are using NMS in the cloud, the NMS SQL Database is also in the cloud and you do not have to allocate space for it.

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 Console*, but for all data generated at most installations with up to 500 dictating providers.

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 Medical Enterprise*'s future speech recognition ability, the server does not collect the data; instead each *Dragon Medical Client* on the network collects data while the healthcare provider 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 Profile Optimizer Server and Profile Optimizer database

The Profile Optimizer Server and Profile Optimizer Database always reside within your network. The *Profile Optimizer Database* often does not require as much space as the *NMS* database. This database contains the scheduling information for ACO and LMO tasks that take place on the *Profile Optimizer Speech Nodes*.

By default the *Profile Optimizer* creates a 10 MB file to contain the scheduling data in the database.

The Profile Optimizer Server stores all of the logs for the Profile Optimizer tasks, 50 log files for the Profile Optimizer Server service. You should reserve between 500 MB and 1 GB of storage space for all the log files the Profile Optimizer Server creates and stores.

Since the total space required for the *Profile Optimizer Server* and the *Profile Optimizer Database* is relatively small, you usually install both of them on the same physical server.

About disk space for server installations and logs

In addition to considering how much space each database requires, you should also be planning for space for each server installation and for logs that the *Profile Optimizer Speech Nodes* generate and store on the *Profile Optimizer Server*.

Each server requires about 20 GB of free space after the installation of the *NMS* and *Profile Optimizer* software.

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 *Profile Optimizer Server* 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 provider every day and have 500 providers, 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 *Profile Optimizer* 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 providers 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 360 | 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 DM360 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 DM360 Network Edition *Installation Guide* or the *Nuance Management Server 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 360 | Network Edition network, you must open particular ports to ensure the free flow of data between servers, clients, and applications. See the diagram and table below for a list of ports in a DM360 Network Edition network. You must open all ports listed in the table on any hardware firewalls to ensure that no firewall blocks transmission of data and brings the network to a halt.

Notes:

- In a DMZ, the NMS server requires access to LDAP/Active Directory to access firewalls and verify user login information.

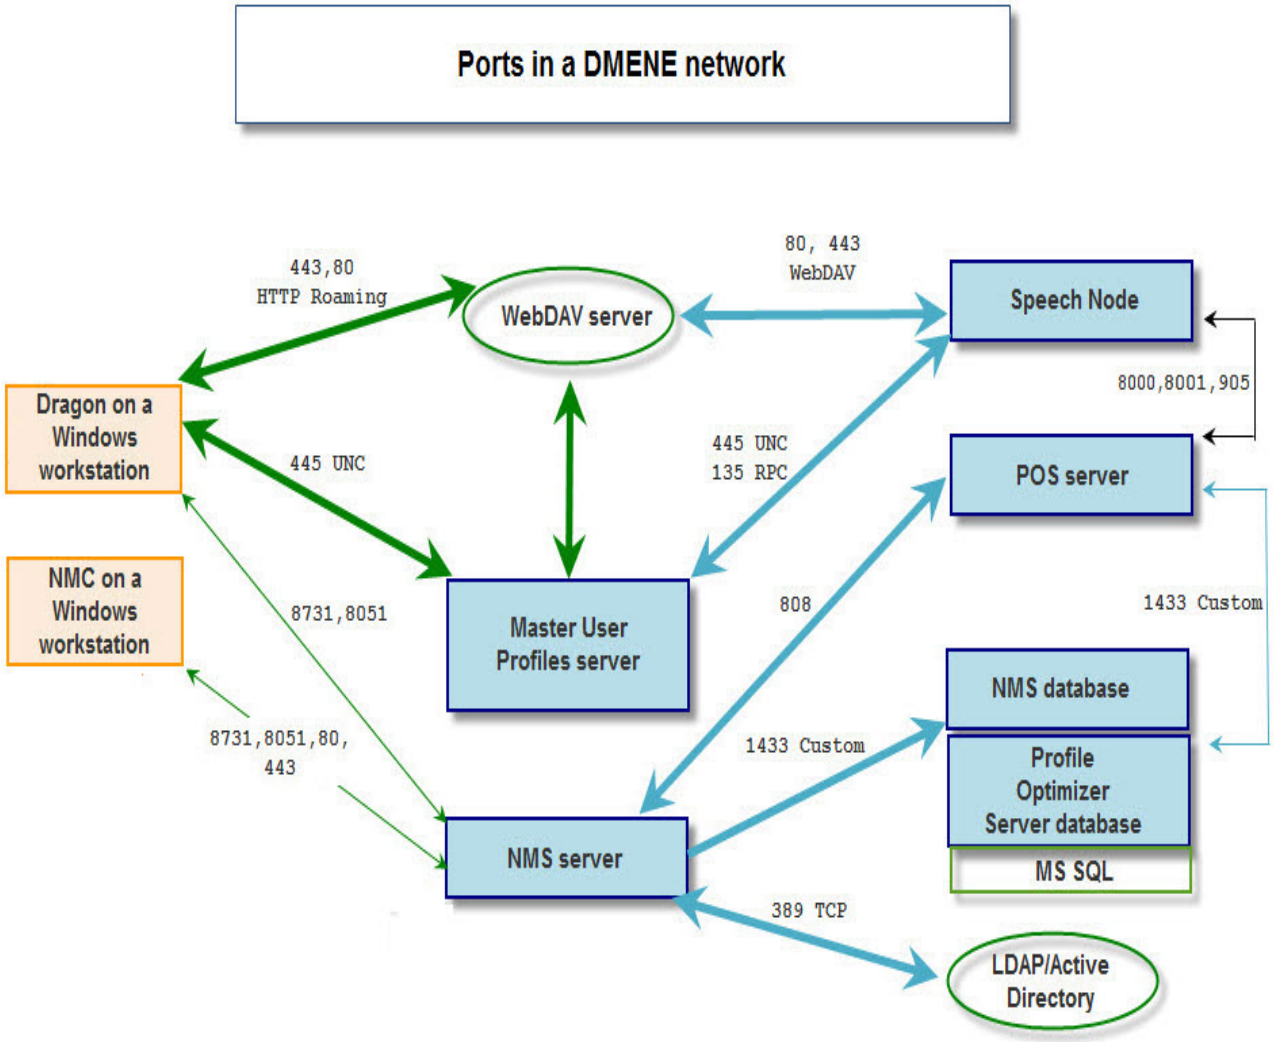

You must open the following ports in the DM360 Network Edition network:

- Ports between workstations where you expect to run the *Nuance Management Console* and the *NMS Server*
- Ports between *NMS Server* and its database if that database is on a separate server
- Ports between *Profile Optimizer Server* and its database if the database is on a separate server
- Ports between *NMS Server* and *Profile Optimizer Server*
- Ports between *Profile Optimizer* and each *Speech Node* machine
- 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

For more information including how to verify access to ports, see 'Test access to NMS ports from a site client machine' in the DM360 Network Edition Installation Guide.

Summary: Ports to open on Clients, Servers, and hardware firewalls

Connection start point	Connection end point	Port to open	Protocol	Description
Dragon on a Windows workstation	WebDAV server	443	HTTP	HTTP Roaming connection to a WebDAV server.
	WebDAV server	80	HTTP	HTTP Roaming connection to a WebDAV server.
	Master User Profiles server	445	UNC	UNC network connection to the Master User Profiles server.
	NMS server	8731	TCP/IP	Connection to the NMS server.
	NMS server	8051	TCP/IP	Connection to the NMS server.
Nuance Management Console on a Windows workstation	NMS server	8731	TCP/IP	Connection to the NMS server.
	NMS server	8051	TCP/IP	Connection to the NMS server.
	NMS server	80	TCP/IP	Connection to the NMS server.
	NMS server	443	TCP/IP	Connection to the NMS server.
WebDAV server	Dragon client	443	HTTP	HTTP Roaming connection to Dragon on a Windows workstation.
	Dragon client	80	HTTP	HTTP Roaming connection to Dragon on a Windows workstation.
	Master User Profiles server	445	TCP/IP	Connection to the Master User Profiles server.
	Speech Node	443	TCP/IP	WebDAV connection to a Speech Node.
	Speech Node	80	TCP/IP	WebDAV connection to a Speech Node.
NMS Server	Dragon client	8731	TCP/IP	Connection to Dragon on a Windows workstation.
	Dragon client	8051	TCP/IP	Connection to Dragon on a Windows workstation.
	Nuance Management Console	443	HTTPS	SSL connection to the Nuance Management Console on a Windows workstation.
	Nuance Management Console	80	HTTP	HTTP connection to the Nuance Management Console on a Windows workstation.
	Master User Profiles server	445	UNC	UNC connection to the Master User Profiles server.
	NMS SQL database	1433		Custom connection to the NMS SQL database.
	POS service	905	TCP/IP	Connection to the POS server.
	Active Directory	389	TCP/IP	LDAP connection to the Active

Connection start point	Connection end point	Port to open	Protocol	Description
	location			Directory.
Profile Optimizer Server (POS server)	Speech Node	8000	TCP/IP	Connection to a Speech Node.
	Speech Node	905	TCP/IP	Connection to a Speech Node.
	NMS server	808	TCP/IP	Connection to the NMS server.
	NMS database POS database with MS SQL	1433	TCP/IP	Custom connection to the NMS database and the POS database with MS SQL.
Speech Node	WebDav server	80	TCP/IP	WebDAV connection to a WebDav server.
	WebDav server	443	TCP/IP	WebDAV connection to a WebDav server.
	Master User Profiles server	445	UNC	UNC connection to the Master User Profiles server.
	Master User Profiles server	135	RPC	Remote Procedure Call (RPC) connection to the Master User Profiles server.
	POS server	8001	TCP/IP	Connection to the POS server.
	POS server	905	TCP/IP	Connection to the POS server.
Master User Profiles Server	WebDAV server	445	TCP/IP	Connection to a WebDAV server.
	Dragon client	445	UNC	UNC connection to Dragon on a Windows workstation.
	Speech Node	445	UNC	UNC connection to a Speech Node.
	Speech Node	135	RPC	Remote Procedure Call (RPC) connection to a Speech Node.
	NMS server	445	UNC	UNC connection to the NMS server.
NMS database	NMS server	1433	TCP/IP	Custom connection to the NMS server. Standard SQL default port - if the NMS server and the NMS server SQL database are on different physical servers.
POS database with MS SQL	POS server	1433		Custom connection to the POS server.
LDAP/Active Directory	NMS server	389	LDAP	Connection to the NMS server or the Call Forwarder.

Chapter 3: Configuring DM360 Network Edition networks that use physical servers

This chapter presents information about different configurations for DM360 Network Edition networks that primarily use physical servers.
This section includes system requirements for DM360 Network Edition network configurations that utilize physical servers. If you plan to set up a DM360 Network Edition network that uses many physical servers, make sure the network meets the system requirements specified in this section.

You can configure your network for Dragon Medical 360 | Network Edition in several ways.

Using physical servers in a DM360 Network Edition network **34**

System requirements for small networks with physical servers **40**

System requirements for medium networks with physical servers **41**

System requirements for large networks with physical servers **43**

System requirements for master user profile computers **44**

Summary: System requirements for DM360 Network Edition networks that use physical servers **47**

Summary: Configuration options for DM360 Network Edition networks that use physical servers **50**

Using physical servers in a DM360 Network Edition network

Configuring a DM360 Network Edition network for up to 100 providers

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

You can have more than one master user profiles directory if you have more than one site. For more information, refer to the *Nuance Management Server (NMS) Administrator Guide*.

If you want to install all *NMS* (either the full NMS or the Call Forwarder) and *Profile Optimizer* components on a single machine, you should not have more than 100 dictating providers. This configuration requires a high powered physical server to run both the *NMS* components and the *Profile Optimizer Servers* alongside a *Profile Optimizer Speech Node* while storing the master user profiles—all on the same machine.

A quad-core processor allows the *NMS* and *Profile Optimizer* tasks and databases to run on the first three cores while the *Profile Optimizer Speech Node* runs acoustic and language model optimizations on the last core. 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.

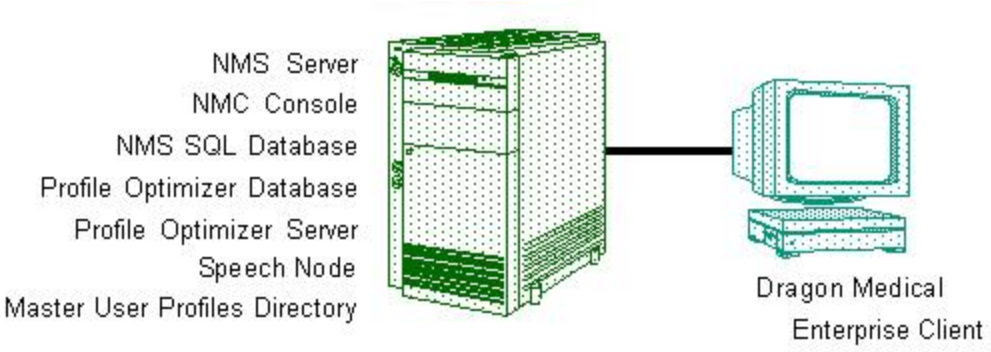

Configuring a DM360 Network Edition network for 100 to 1000 providers

If you have 101 to 1,000 individuals dictating with *Dragon Medical Clients*, you need to begin to scale up your Dragon Medical 360 | Network Edition system to manage the greater workload. One way to expand your system is to install various components of it on separate machines.

Although you have two databases if you are installing the NMS in your network, one for the *NMS* and one for the *Profile Optimizer*, both are small enough to install on one machine with a single

SQL Server license. Whenever it is possible, Nuance recommends installing the databases on the same machine. Nuance requires a dual-core physical server for the*NMS Server*, so that it runs more efficiently, because the server might need to run compression/encryption on up to 600 MB of data for a single user profile to support data collection.

The most optimal distributed configuration installs the *NMS* and *Profile Optimizer* components 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 Console or NMS Call Forwarder
    - *NMS SQL Database*
    - *Profile Optimizer Server*
    - *Profile Optimizer Database*
    - Master user profiles directory
- Second server software components:
    - Up to 4 *Profile Optimizer Speech Nodes*, each on its own virtual machine.

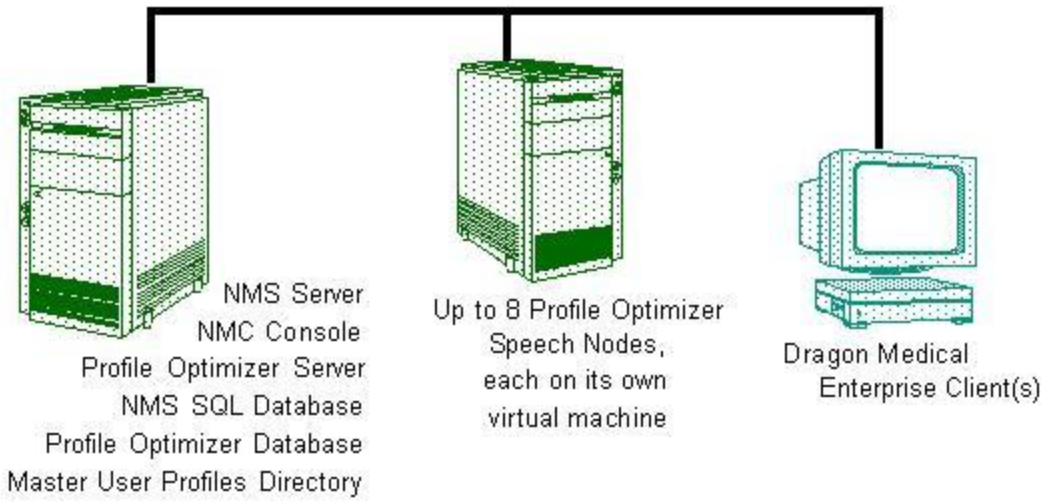

Configuring a DM360 Network Edition network for more than 1000 providers

If you have a larger organization with more than 1,000 dictating healthcare providers, 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:
    - One *NMS Server* or Call Forwarder for every 1,000 providers
    - A *Nuance Management Console* for each *NMS Server*

- A single Database server with these components:
    - *Profile Optimizer Server*
    - *NMS SQL Database*
    - *Profile Optimizer Database*
    - Master User Profiles directory
- At least one eight-core server with these components:
    - Up to 4 *Profile Optimizer Speech Nodes*, one for every 125 providers

Note:

You can have more than subdirectory inside the master user profiles directory, one for each site. For more information about setting up sites, refer to the *Nuance Management Server Administrator Guide*.

The illustration that follows shows the most basic three server configuration that you would use with 1,000 providers, the minimum number to require a large configuration.

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.

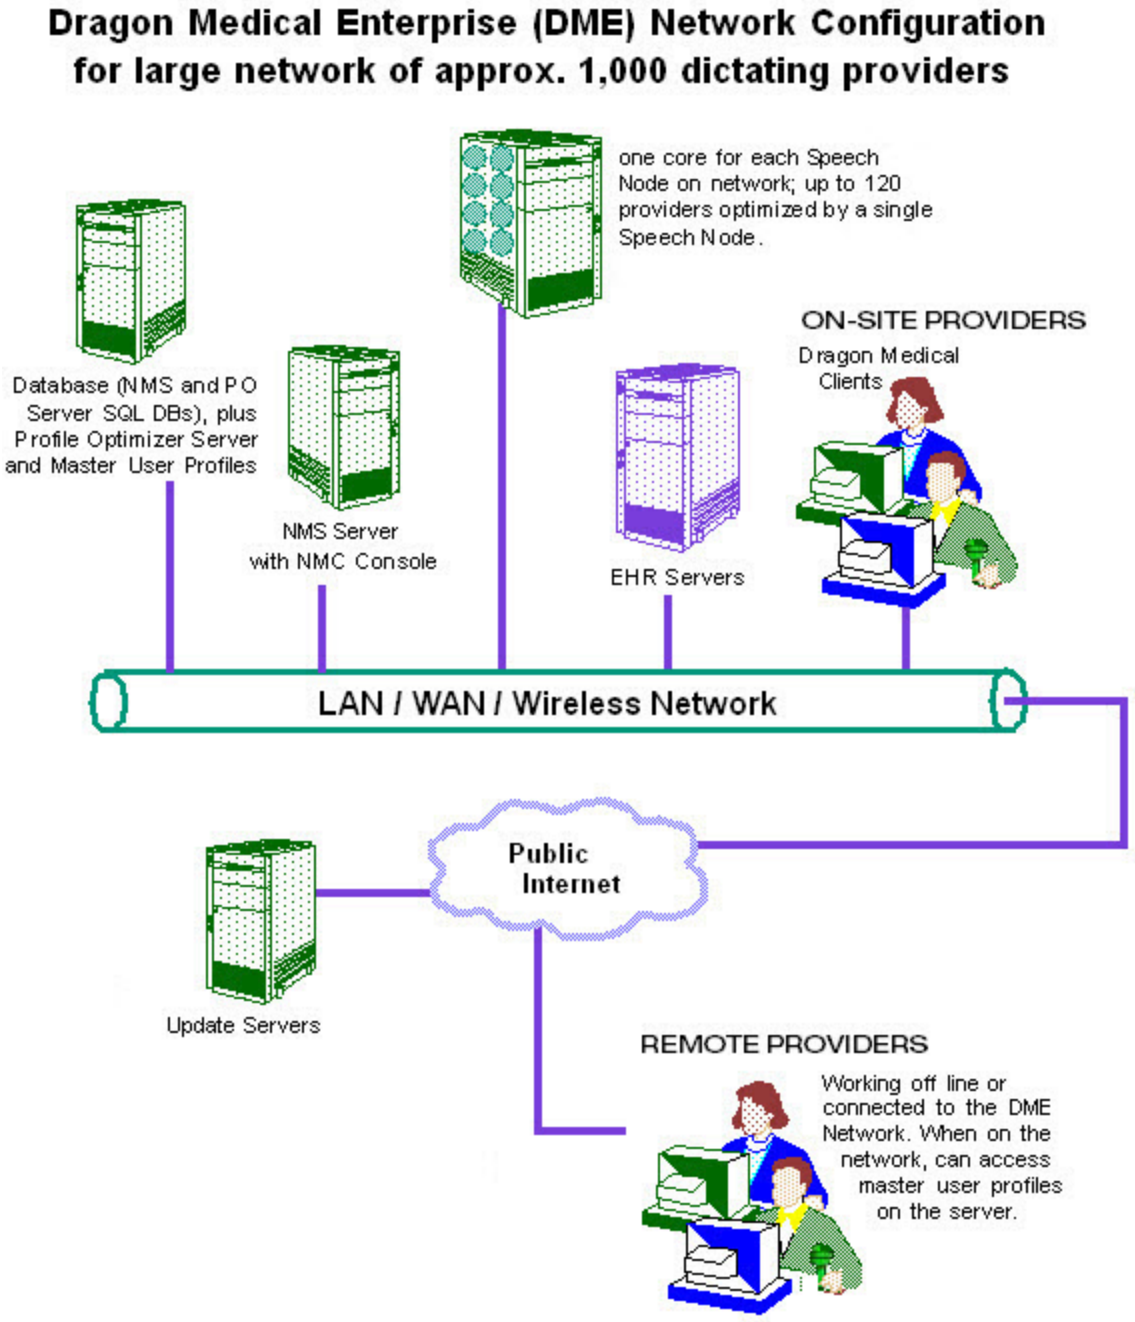

Note: Speech node servers need two cores for every Speech Node, for example quad-core for two Speech Nodes.

The *NMS Server* must be dual-core, because the *NMS Server* takes advantage of the extra core to provide management of your Dragon Medical 360 | 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 DM360 Network Edition network configuration for 4,000 healthcare providers. 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*.

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

The next illustration shows the master user profiles stored on a remote SAN storage device. In this scenario, the *Profile Optimizer Server* receives requests from the *Speech Nodes* that want to access a profile and remotes those calls to the Windows server that the SAN storage device is connected to through a Host Bus Adaptor card on a 10/100/1000 Mbps Ethernet.

Remote Master User Profiles on SAN storage device and Profile Optimizer remotes calls to its server

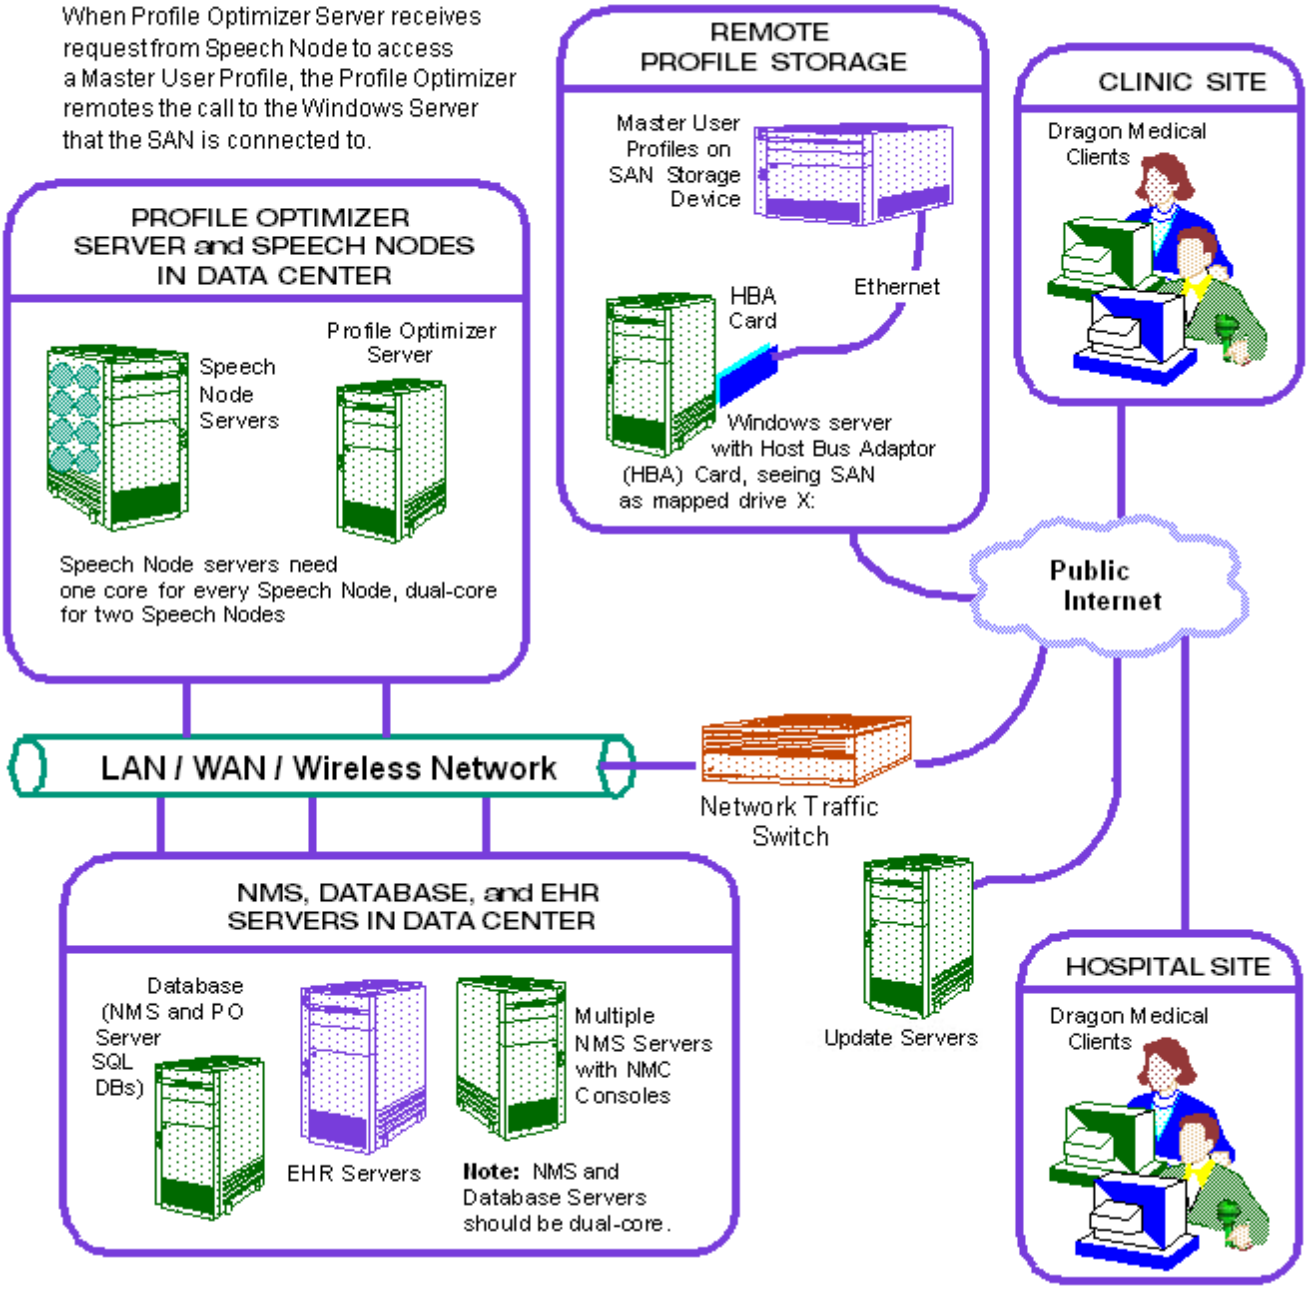

System requirements for small networks with physical servers

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

- **Single Quad-Core Physical Server**
  - **One Quad-Core Server** for *NMS Server* & *Nuance Management Console* or NMS Call Forwarder, *NMS SQL* and *Profile Optimizer Databases* & *Profile Optimizer Server/Scheduler* 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 2003 SP2 or R2, 32- or 64-bit version.
    - Windows Server 2008 or with SP2 (optional)
    - Windows Server 2008 R2
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Microsoft .NET**:
    - Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required for the server.
    - Microsoft .NET Framework 4.0 client or full version for the Nuance Management Console.
- **Minimum RAM**: 4 GB recommended
- **Core Application Disk Storage**: 40 GB for the NMS Server and Profile Optimizer Server
- **Maximum Number of Speech Nodes**:
  - If using only 1 physical server, it needs to be a Virtual Server Host. NMS and POS 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 providers
- **Master User Profile Storage**:
  - Master user profiles directory on its own RAID array
  - Master user profiles on their own totally separate server or workstation
  - See System requirements for Master User Profiles on a Profile Optimizer Server for details.

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 providers.

- **Two Physical Servers, One Quad-Core and One Eight-Core**
  - **First Server (Quad-Core):** *NMS Server* & *NMC Console* or NMS Call Forwarder, *NMS SQL* and *Profile Optimizer Databases*, and *Profile Optimizer Scheduler/Server*
  - **Second Server (Eight-Core):** Up to 4 *Profile Optimizer Speech Nodes* (one for each two cores)
  - Master User Profiles Directory on first server or on a separate third machine (server not required, but recommended)
- **First server requirements**:
  - **Processor**: Dual-Core 2 GHz CPU
  - **Operating System**:
    - Windows Server 2003 SP2 or with R2, 32- or 64-bit version.
    - Windows Server 2008 with SP2 (optional) or R2
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Microsoft .NET**:
    - Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required.
    - Microsoft .NET Framework 4.0 client or full version for the Nuance Management Console.
  - **Minimum RAM**: 4 GB RAM for NMS, 4 GB RAM for Profile Optimizer
  - **Approximate Disk Storage**:
    - 20 GB for the NMS Server
    - 20 GB for the Profile Optimizer Server
  - **Master User Profile Storage**:
    - Master user profiles directory on its own RAID array
    - See System requirements for Master User Profiles on a Profile Optimizer Server for details.
- **Second server requirements**:
  - **Processor**: Eight-Core 2 GHz CPU
  - **Operating System**:
    - Windows XP SP3*
    - Windows Vista SP2
    - Windows 7
    - Windows Server 2003 SP2, or with R2, 32-bit or 64-bit.

- 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 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required.
- **Minimum RAM**: 16 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 providers

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

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 providers.

- **Minimum of Three Physical Servers, all Eight-Core**
  - **One *NMS Server*** or Call Forwarder for every 1,000 providers
  - **One Database Server** for *NMS SQL* (if used) and *Profile Optimizer Databases*, plus the *Profile Optimizer Scheduler/Server* and Master User Profiles Directory
  - **One Eight-Core Server** for every 4 *Profile Optimizer Speech Nodes*, allowing one node for every 1,000 providers
- **First server requirements**:
  - **Processor**: Dual-Core 2 GHz CPU
  - **Operating System**:
    - Windows Server 2003 SP2 or with R2, 32- or 64-bit version
    - Windows Server 2008 with SP 2 (optional) or with R2
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Microsoft .NET**:
    - Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required for the server.
    - Microsoft .NET Framework 4.0 client or full version for the Nuance Management Console.
  - **Minimum RAM**: 4 GB RAM
  - **Approximate Disk Storage**: 20 GB for each *NMS Server* instance
- **Second server requirements**:
  - **Processor**: Dual-Core 2 GHz CPU
  - **Operating System**:
    - Windows Server 2008 with SP 2 (optional) or with R2
    - Windows Server 2003 SP2 or with R2, 32- or 64-bit version.
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Minimum RAM**: 4 GB RAM
  - **Minimum Disk Space**: 20 GB for the *Profile Optimizer Server*
    - 20 GB for the Profile Optimizer Server
  - **Master User Profile Storage**:
    - Master user profiles directory on its own RAID array
    - See System requirements for Master User Profiles on a Profile Optimizer Server for details.
- **Third server requirements**:

- **Processor**: Eight-Core 2 GHz CPU
- **Operating System**:
    - Windows XP SP3*
    - Windows Vista
    - Windows 7
    - Windows Server 2003 SP2 or with R2, 32-bit or 64-bit version
    - 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 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required.
- **Minimum RAM**: 16 GB RAM
- **Approximate Disk Storage**: 20 GB for each *Profile Optimizer Speech Node*
- **Minimum and maximum number of providers**: Up to 1,000 dictating providers 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

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

System requirements for master user profile computers

Any time after you have installed the *NMS Server* and started the *NMS Server* service (even before you have the *Profile Optimizer Server* installed or running), 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:

1 GB RAM for Windows XP Professional* with SP3 or Windows Vista, 2 GB RAM for Windows Server 2003 SP2 or Windows 7, 32-bit and 64-bit. *Recommended:* 4 GB RAM for Windows 7, 64-bit, 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 2003 SP2
- Windows XP Professional with SP3*
- Windows Server 2008
- Windows Vista
- Windows 7
- 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 3.5 SP1 installed.

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

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 3.5 SP1 installed.

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 and having the *Profile Optimizer* access the profiles by making remote calls to that server. 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 360 | Network Edition server (usually the *Profile Optimizer Server* or another network server running Windows 2003 or 2008 Server with .NET Framework 3.5 SP1 and .NET Framework 4.0 installed) 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:).

The *Profile Optimizer* remotes calls from the *Speech Nodes* to access the master user profiles to the server that the SAN is connected to.

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 or locate them on the same machine with any of the following components: *NMS Server*, *NMS SQL Database*, *Profile Optimizer Server*, or *Profile Optimizer 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 Medical Client* workstations would access the profiles on a web server through a URL starting with **http://**.
- *Dragon Medical Client* workstations would access the profiles on a secure (SSL) web server through a URL starting with **https://**.
- *NMS Server* and *NMC Console* 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 DM360 Network Edition *Installation Guide*.

Summary: System requirements for DM360 Network Edition networks that use physical servers

	NMS SQL Database	Profile Optimizer Database	NMS Server (or Call Forwarder) and NMC Console	Profile Optimizer Server	Speech Nodes	Master User Profiles Host	Dragon Medical Client
Computer Type	Server	Server	Server	Server	Workstation or Server	Workstation or Server	Workstation
Processor	Xenon 3050 or equiv-alent	Xenon 3050 or equiv-alent	Xenon 3050 or equivalent	Xenon 3050 or equiv-alent	One core per node	At least Pentium IV	Pentium IV or AMD Athlon 64
CPU Speed	2 GHz	2 GHz	2 GHz	2 GHz	2 GHz	2 GHz	2.4 GHz (AMD 1 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	20 GB	20 GB	20 GB	Based on calculation	2 GB; Recommended: 2.5 GB
Minimum RAM	4 GB	4 GB	4 GB	2 GB	2 GB per node, up to 16 GB on an eight-core server	4 GB	512 MB on Win-dows XP*, 1 GB on Windows Vista, 2 GB on Windows 7, 32- & 64-bit; Recommended: 4 GB on Windows 7, 64-bit

	NMS SQL Database	Profile Optimizer Database	NMS Server (or Call Forwarder) and NMC Console	Profile Optimizer Server	Speech Nodes	Master User Profiles Host	Dragon Medical Client
Minimum L2 cache	2 MB	2 MB	2 MB	2 MB	2 MB	2 MB	512 KB; Recommended: 1 MB L2
Other Hardware items	---	---	---	---	---	---	Sound card recording at 16 bit 11 KHz and speakers for playback; Microphone; DVD reader
Windows Operating System	Windows Server 2008 or 2003 SP2, Windows Server 2012, and Windows Server 2012 R2	Windows Server 2008 or 2003 SP2, Windows Server 2012, and Windows Server 2012 R2	Windows Server 2008, 2008 R2, 2003 SP2, or 2003 R2 (32- or 64-bit), Windows Server 2012, and Windows Server 2012 R2	Windows Server 2008, 2008 R2, 2003 SP2, or 2003 R2 (32- or 64-bit), Windows Server 2012, and Windows Server 2012 R2	Windows XP* SP3, Windows Vista SP2, or Windows 7 workstation; Windows Server 2008 or 2003 SP2, Windows Server 2012, and Windows Server 2012 R2	Windows XP* SP2 or SP3, Windows Vista, or Windows 7 workstation; Windows Server 2008 or 2003 SP2; RAID Array may be required (see note below), Windows Server 2012, and Windows Server 2012 R2	32-bit and 64-bit operating systems detailed below table.
SQL Server	2008	2008	---	---	---	---	---
.NET Framework	Version 3.5 SP1 Required	Version 3.5 SP1 Required	Version 3.5 SP1 and Version 4.0 Client or full version Required	Version 3.5 SP1 and Version 4.0 Client or full version Required	Version 3.5 SP1 and Version 4.0 Client or full version Required	Version 3.5 SP1 Required	---
Internet Information Services (IIS)	---	---	Required	---	---	---	----
In Windows	Required	Required	Required	Not	Not required	Required	Required

	NMS SQL Database	Profile Optimizer Database	NMS Server (or Call Forwarder) and NMC Console	Profile Optimizer Server	Speech Nodes	Master User Profiles Host	Dragon Medical Client
Domain				required			
Windows Installer 3.1 or later	---	---	---	Required	Required	---	---
Dragon SDK Client	---	---	---	---	Required	---	---
Internet Browser	---	---	Internet Explorer 8, 9 and 10, Chrome, Firefox, Safari or any browser with support for WPF xbap technology	---	---	---	Internet Explorer 6, 7, or 8

Notes:

- *Dragon Medical Enterprise Client* software can run on several 32-bit and 64-bit operating systems. See the system requirements section of this chapter for a complete list.
- Whenever you store the master user profiles on the same machine where the *Profile Optimizer Server* resides, you must store those profiles on a RAID array.

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

Summary: Configuration options for DM360 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 Providers	10 - 100	101 - 1,000	1,000 or more
Minimum Physical Servers	1 Server	2 Servers	3 Servers: One Dual-Core NMS Server & one Eight-Core Speech Node Server for every 1,000 providers (one Speech Node server allows for 4 Speech Nodes); One Database Server
NMS Server			
Software	NMS Server or Call Forwarder Profile Optimizer Server NMS and Profile Optimizer Databases Master User Profiles Folder 1 Speech Node	NMS Server or Call Forwarder Profile Optimizer Server NMS and Profile Optimizer Databases Master User Profiles Folder	NMS Server or Call Forwarder
Hardware	Quad-Core with minimum 4 GB RAM Hard Disk = 40 GB + 1 GB per provider	Dual-Core with minimum 4 GB RAM Hard Disk = 40 GB + 1 GB per provider	Dual-Core with minimum 4 GB RAM Hard Disk = 40 GB + 1 GB per provider
Speech Node Server			
Software	---	Up to 8 Speech Nodes on VMs	Up to 8 Speech Nodes on VMs
Hardware	---	Eight-Core with a minimum of 8 GB of RAM. For Windows Vista, Windows 7, Windows Server 2008, Windows Server 2012, and Windows Server 2012 R2, the recommendation is 16 GB of RAM. Hard Disk = 20 GB per node	Eight-Core with a minimum of 8 GB of RAM. For Windows Vista, Windows 7, Windows Server 2008, Windows Server 2012, and Windows Server 2012 R2, the recommendation is 16 GB of RAM. Hard Disk = 20 GB per node
Database Server			

	Size and type of network configuration		
	Small: One server	**Medium: Two servers**	**Large: Three or more servers**
Software	---	---	Profile Optimizer Server NMS and Profile Optimizer Databases Master User Profiles Folder
Hardware	---	---	Dual-Core with minimum 4 GB RAM Hard Disk = 20 GB per node + 1 GB per provider
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 *System requirements for master user profile computers* on page 44. Hard Disk = 20 GB + 1 GB per provider

Chapter 4: Configuring DM360 Network Edition networks that use virtual servers

This chapter presents information about different configurations for Dragon Medical 360 | Network Edition networks that primarily use virtual servers.
This section includes system requirements for DM360 Network Edition network configurations that utilize virtual servers. If you plan to set up a DM360 Network Edition network that uses many virtual servers, make sure the network meets the system requirements specified in this section.

You can configure your network for DM360 Network Edition in several ways.

Using virtual servers in a DM360 Network Edition network 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 57

Using virtual servers in a DM360 Network Edition network

To use virtual machines in your Dragon Medical 360 | 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 or NMS Call Forwarder, the Profile Optimizer, and databases

You can choose to generate medium or larger configurations using virtual machines for the *NMS Server*, *Profile Optimizer Server*, and/or their 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 one full core for each *Speech Node*. On a dual-core machine, you can have two *Speech Nodes* on separate virtual machines; on a quad-core machine, you can have three or four *Speech Nodes* each on its own virtual machine.

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 providers.

- **Single Quad-Core physical or virtual server**
  - **One Quad-Core Server** for NMS Server & Nuance Management Console or NMS Call Forwarder, NMS SQL and Profile Optimizer Databases & Profile Optimizer Server/Scheduler, 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 2003 SP2 or R2, 32- or 64-bit version
    - Windows Server 2008 or with SP2 (optional), 32- or 64-bit version
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Microsoft .NET**:
    - Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required for the server.
    - Microsoft .NET Framework 4.0 client or full version for the Nuance Management Console.
  - **Minimum RAM**: 4 GB is recommended
  - **Core Application Disk Storage**: 60 GB
    40 GB for the Operating system and the NMS Server and Profile Optimizer Server Services
  - **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 providers
  - **Master User Profile Storage**: 1.1 GB per provider/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 providers.

- **Minimum of two virtual servers**
  - **First Server (Dual vCPU)**: NMS Server & Nuance Management Console or NMS Call Forwarder, NMS SQL and Profile Optimizer Databases, and Profile Optimizer Scheduler/Server
  - **Additional Servers (1 per 125 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 2003 SP2 or R2, 32 bit or 64 bit
    - Windows Server 2008 or R2, 32 bit or 64 bit
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Microsoft .NET**:
    - Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required.
    - Microsoft .NET Framework 4.0 client or full version for the Nuance Management Console.
  - **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 125 Users) requirements**
  - **Processor**: 2 vCPU
  - **Operating System**:
    - Windows XP SP3, 32-bit or 64-bit*
    - Windows Vista SP2, 32-bit or 64-bit
    - Windows 7, 32-bit or 64-bit
    - Windows Server 2003 SP2 or R2, 32-bit or 64-bit
    - Windows Server 2008, 32-bit or 64-bit
    - Windows Server 2008 R2: On all machines where Speech Nodes will 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%20R2%20SP1/sp2/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

  - Windows Server 2012
  - Windows Server 2012 R2
- **Microsoft .NET**: Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required
- **Minimum RAM**: 1 GB RAM for Windows XP*, 2 GB RAM for Windows Vista SP2, Windows 7, Windows Server 2003, Windows Server 2008, Windows Server 2012, and Windows Server 2012 R2
- **Hard Disk Space**: 20 GB for Windows XP*, 30 GB for Windows Vista SP2, Windows 7, Windows Server 2003, Windows Server 2008, Windows Server 2012, and Windows Server 2012 R2
- **Minimum and Maximum Number of Users**: 101 to 1000 dictating providers
- **Master User Profile Storage**: 1.1 GB per provider/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.

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

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 providers.

- **Minimum of three virtual servers**
  - **One *NMS Server*** or Call Forwarder for every 1,000 providers
  - **One Database Server** for *NMS SQL* and *Profile Optimizer Databases* or a customer SQL cluster
  - **One Speech Node**: one node for every 125 providers
- **First server: NMS Server hardware requirements and limitations**
  - **Processor**: 2 vCPU
  - **Operating System**:
    - Windows Server 2003 SP2 or R2, 32- or 64-bit version
    - Windows Server 2008, R2, or SP2 (optional), 32- or 64-bit version
    - Windows Server 2012
    - Windows Server 2012 R2
  - **Microsoft .NET**:
    - Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 4.0 Client or full version required for the server.
    - Microsoft .NET Framework 4.0 client or full version for the Nuance Management Console.
  - **Minimum RAM**: 4 GB RAM
  - **Approximate Disk Storage**: 40 GB
- **Second server: Database server hardware requirements and limitations**
  - **Processor**: 2 vCPU or Existing SQL Cluster
  - **Operating System**:
    - Windows Server 2003 SP2 or R2, 32- or 64-bit version
    - 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
  - **Database**: MS SQL Server 2008 Standard or higher
- **Third server: Speech Nodes (Speech Nodes 1 per 125 Users)**
  - **Processor**: 2 vCPU
  - **Operating System**:
    - Windows XP SP3, 32-bit or 64-bit*
    - Windows Vista, 32-bit or 64-bit

- Windows 7, 32-bit or 64-bit
- Windows Server 2003 SP2 or R2, 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 3.5 SP1
  - Microsoft .NET Framework 4.0 Client or full version required
- **Minimum RAM**: 1 GB RAM for Windows XP*, 2 GB RAM for Windows Vista, Windows 7, Windows Server 2003, Windows Server 2008, Windows Server 2012, and Windows Server 2012 R2
- **Hard Disk Space**: 20 GB for Windows XP*, 30 GB for Windows Vista, Windows 7, Windows Server 2003, Windows Server 2008, Windows Server 2012, and Windows Server 2012 R2
- **Minimum and Maximum Number of Users**: Over 1000 dictating providers
- **Master User Profile Storage**: 1.1 GB per provider/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.

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

Chapter 5: System requirements for the Dragon client and user profiles on a Profile Optimizer Server

This chapter presents the system requirements for the DM360 Network Edition Dragon client and for storing Master User Profiles on a Profile Optimizer Server. Use the information in this chapter as guidelines when you set up your DM360 Network Edition network.

System requirements for the Dragon Medical Client **60**

System requirements for Master User Profiles on a Profile Optimizer Server **62**

System requirements for the Dragon Medical Client

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

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

DVD-ROM: drive required for installation

Free hard disk space: 5 GB

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

Microphone: Nuance-approved microphone (included in purchase)
For details on Bluetooth microphones, recorders, Tablet PCs, and other hardware, please go to support.nuance.com/compatibility/.

- Standard Headsets
- USB microphones
- PowerMic 2 non-scanner
- PowerMic 2 scanner
- Plantronics Calisto Bluetooth microphones
- Digital recorders
- Enhanced Bluetooth (Calisto II) microphones
- Special enrollment for array microphones
- an iOS or Android device as a microphone using the Dragon Remote Microphone app
- Other Microphones listed in support.nuance.com/compatibility/

Operating system:

- 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, and 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

Processor Cache: 2 MB

RAM:

- 2 GB RAM for Microsoft® Windows XP®* and Microsoft® Windows Vista®.
- 4 GB for Microsoft® Windows® 7, 32-bit and 64-bit, Microsoft® Windows® 8, 32-bit and 64-bit, Windows Server 2008 64-bit, Windows Server 2012 64-bit, and Windows Server 2012 64-bit R2.

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

Supported non-EHR Applications:

- DragonPad
- WordPad
- NotePad
- Microsoft® Word 2007, 2010 (32 & 64 bit), 2013 (32 & 64 bit)
- Microsoft® Outlook® 2007, 2010, 2013
- Microsoft® Excel® 2007, 2010, 2013
- WordPerfect® x5, x6
- Apache OpenOffice Writer 3.4
- Open Office Writer v3.1, 3.2
- Internet Explorer 8, 9, 10
- Rich Internet Application IE9
- 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

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

System requirements for Master User Profiles on a Profile Optimizer Server

If you decide to store the master user profiles on the same machine as the *Profile Optimizer*, you should be sure that you install:

- Master user profiles are on their own RAID array.
- .NET Framework 4.0 and .NET Framework 3.5 SP1 on the server where the master user profiles are stored.

Notes: RAID array setup notes:

- If you can access the drive of the RAID array from the **MyComputer** window of the server without a mapped network connection, then you should treat the master user profiles directory as located on the server and install the required software.
- If you must access the drive of the RAID array over a mapped or other network connection, treat the array as a separate machine.

Chapter 6: Network bandwidth requirements on DM360 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 Requirements

Naturally, *Dragon Medical Clients* are going to operate more fluidly in a faster network. And the more healthcare providers 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 providers 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 provider is going to experience remains reasonable?

You want to avoid setting up an underpowered network—where provider 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 Medical 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 Medical 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 provider select the **If you know your network is slow, enable this option to work offline** check box in the **Dragon Login** dialog box (see next illustration).

The *NMS Server* will still validate the provider's license, but not communicate with the *Master User Profile Server* while the provider is dictating. The provider 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 provider 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 Medical Client* workstation.

Note: This setting differs from setting up the provider to work in Disconnected mode (by checking the **Allow disconnected mode** check box in the Site setup). The provider who is working in Disconnected mode has up to 90 days of access to the client before being forced to log in over the network. Providers working in Disconnected mode see similarly improved performance, because they are not interacting with the network.

Considerations for Network Bandwidth in Citrix Environments with vSync

The CPU load that vSync adds is negligible. vSync uses approximately 10 MB of memory per provider dictation session.

Estimating Load and Traffic Generated in Citrix Environment Using vSync Feature

Another factor that might influence the function of your DM360 Network Edition *Network* is setting the network up with a Citrix Server and using the vSync feature of DM360 Network Edition. vSync introduces a small delay that you can add to the delays stated earlier.

Load Created by 40 Providers Logging in on DM360 Network Edition Network with Dragon Using the vSync Feature

Load	Server idle	One provider dictating	40 providers with EHR and vSync	40 providers 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 providers utilize 2.7% of the server's CPU and that each provider load is approximately 10 MB of memory (see the **vSync Load** column to the far right).

Other load information on *Dragon Medical Clients* with vSync:

- With vSync loaded, but *Dragon Medical Client* not active, no significant/measurable increase in CPU usage occurs.
- With vSync loaded, but *Dragon Medical Client* not active, an increase in memory usage of approximately 4 MB per dictating provider occurs.

- With vSync loaded and *Dragon Medical Client* active, an additional increase in memory usage of approximately 5 MB per dictating provider occurs (for a total increase of approximately 9 MB per provider resulting from both vSync and dictating in Dragon).

The traffic created by vSync feature is summarized below:

- Sites using vSync with *Dragon Medical Clients* would see about 1 KB of traffic for every 100 words dictated. This amount of traffic is similar to traffic generated by instant messaging.

For hardware and software specifications of the testing environment that produced these numbers, see Considerations for Low Network Bandwidth Configurations.

Resolving Communication between EHR in Citrix and Full Text Control

When the provider connects to the Citrix server, vSync starts automatically in the Citrix session.

When a provider begins dictating into *Dragon* in a Citrix environment or into a supported Citrix-published EHR, the Full Text Control indicator (see illustration on next page) on the *DragonBar* turns green and remains green as long as the *Dragon Medical Client* is communicating successfully with the vSync module. If the Full Text Control color changes from green to gray, it means one of the following conditions exists:

- The Citrix published EHR is using an unsupported Edit Control.
- vSync did not send back EHR Edit Control information within the allotted time (2000 ms).
- vSync is not running for some reason.

If vSync does not respond within the 2000 ms allotted, the reason might not be network performance, but could instead be one of the following:

- Your Citrix server has too many sessions running. If Citrix servers are overloaded, that will slow down response of vSync's communication with the EHR.
- Your EHR is slow for some other reason.
- Your EHR database might be slow.
- vSync might not respond to a *Dragon Medical Client* request because communication on the server between the EHR and vSync is too slow.
- Another application in your network is producing excess traffic.

If during dictation the *Dragon Medical Client* does not receive response from vSync within the 2000 ms allotted, the client times out; vSync might generate a warning to the client and the Full Text Control on the *DragonBar* might turn from green to gray, indicating that NatText is handling the words just spoken. Under these conditions, dictating with *Dragon* will still place recognized text in your EHR application, but Full Text Control capabilities will not be possible as long as the indicator is gray.

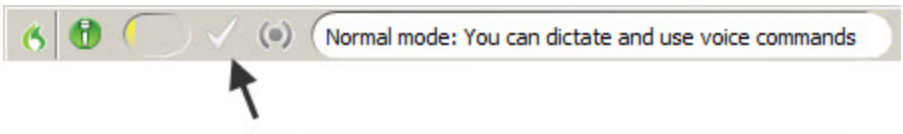

When vSync does not respond, the check mark becomes gray. If the check mark changes between green and gray, this may indicate that a network issue exists.

You can take these actions in a situation where vSync seems to be unresponsive (and the problem is not likely to be a network issue):

1. Turn off vSync for Citrix traffic on the client. Once you turn it off, vSync remains off until the provider exits and restarts/logs in to *Dragon Medical Client* again.
2. Close and reopen the user account.
3. Close the EHR window you are working in and reopen it.
4. Restart the Dragon Medical Client.
5. Restart the EHR.

If the Full Text control flashes between green and gray, that indicates a network issue. You can then proceed as explained in *Considerations for Low Network Bandwidth Configurations* on page 64. If you continue to have problems, contact your Nuance representative.

Estimating Dragon Response Times When Provider 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 Medical Clients* in *DM360 Network Edition Network*.

Measurements of Dragon Medical Client Response Times on Network

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

Total Active Providers	Minimum Login time	Maximum Login time	Mean Login time	Range	Mode	Median
No providers 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 providers simultaneously taking actions that require accessing the *NMS Server* over the network.

Although you can configure up to 1,000 providers on each *NMS Server*, typically fewer than 25% of those professionals would be actually dictating at one time and only a percentage of those providers (perhaps 3% to 5% of your total provider 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.

For hardware and software specifications of the testing environment that produced these numbers, see Test Environment.

Approximate Sizes of Data Transferred Over DM360 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 *DM360 Network Edition Network* is the amount of data transferred over the *DM360 Network Edition Network* when the healthcare provider 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 Provider for Each on 100 Mbps Network

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
Provider creates new user profile	Profile created on client and stored on *Master User Profile Server*. Occurs only once for any particular provider.	70 MB +20 MB for additional vocabulary +60 MB for additional dictation source	Dragon client	Master User Profile Server	Access network at user open/close only
Provider saves profile during dictation	Small profile changes copied from client to *Master User Profile Server*.	Up to 15 MB	Dragon client	Master User Profile Server	Saving user also saves local acoustics

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
Provider logs in to the NMS server through the client for the first time	User profile copied from *Master User Profile Server* to client; settings sent from *NMS Server* to client.	70 MB	Master User Profile Serve	Dragon client	
		10 KB	NMS Server	Dragon client	
Provider logs in to the client after the first time (after initial log in, but not first time after ACO)	(Client is accessing user profile from cache) *NMS Server* passes settings to client.	10 KB	NMS Server	Dragon client	
Provider dictates in Disconnected mode, then reconnects to network	Data from Disconnected mode sessions and changed settings copied from local cache to *Master User Profile Server*.	Less than 1 MB + size of the dictation audio files (1.3 MB/min)	Dragon client	Master User Profile Server	Allow Disconnected Mode
Provider logs in for First time after an ACO completes	An updated user profile copied from *Master User Profile Server* to client.	70 MB	Master User Profile Server	Dragon client	
Provider dictates and makes corrections	Incremental changes to user profile copied to *Master User Profile Server*.	1.3 MB per minute of audio	Dragon client	Master User Profile Server	Conserve server archive size
Provider logs out of Dragon	User profile changes copied to *Master User Profile*	1 MB + 1.3 MB/min of DRA (audio) files	Dragon client	Master User Profile Server	

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
	Server				
Provider carries out Training	User profile changes copied to *Master User Profile Server*.	15 MB	Dragon client	Master User Profile Server	
ACO	*Profile Optimizer* starts ACO on *Speech Node*.	Up to 500 MB	Master User Profile Server	Speech Node	ACO Scheduled
	After ACO completes, *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	PO Server Database	Copy Dragon log to the network
	NMS Server creates foundation service log file.	10 MB	NMS Server	NMS Server Database	
LMO—Files copied to *Speech Node* only once and integrated into multiple profiles.	*Profile Optimizer Server* initiates LMO process on *Speech Node*.	Less than 1 MB, depends on vocabulary changes; same single copy of the data used in all profiles, so effect on network is insignificant.	NMS Server	Speech Node	

Client activity	Network activity	Data transfer size	Transfer source	Transfer destination	Dragon setting
Medication or Command Updates—Files copied to Speech Node only once and integrated into multiple profiles. Because only copied once and files are under 1 MB, generates minimal traffic.	*Profile Optimizer Server* integrates this data when it determines the integration is necessary and performs an LMO-like process to integrate into *Master User Profile Server*; the process runs on the *Speech Node*, but data is not transferred to the client until the next time the provider logs in.	Less than 1 MB, depends on number of medications/commands; same single copy of the data used in all profiles, so effect on network is insignificant.	Update Server	NMS Server	
			NMS Server	PO Server	
			PO Server	Speech Node	
			Speech Node	Profile Server	

Process Producing Greatest Network Load

Network load is affected most when providers are in the process of creating a new user profile. During this process, *DM360 Network Edition* transfers 70 MB of data from the client to the *Master User Profile Server*. *DM360 Network Edition* transfers additional data for additional vocabularies and dictation sources. Usually a provider has only one vocabulary, such as **Cardiology**, based on a specialty. At most a provider might have two vocabularies.

Also, *DM360 Network Edition* transfers additional data if the provider has more than one dictation source/device, whether that means multiples types of microphones, multiple types of recorders, or both. Providers likely have a maximum of two types of microphones and/or two types of recorders and many have only one type.

Since each provider 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 Provider's User Profile

For each workstation where a provider dictates, *DM360 Network Edition* must download the provider's user profile to that machine *only the first time* the physician uses *Dragon* on that workstation. That initial download is 70 MB of data; however, that download occurs only once for a provider per workstation that the provider dictates on. For providers who work on the same one or two workstations most of the time, this action occurs only on first using the workstation and, later, after an ACO 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 Medical 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 *Profiles Server* is newer is immediately after an Acoustic Model Optimization (ACO) process has occurred. After the ACO occurs, the client must download the master user profile for the provider during the next login to the *NMS Server*. Since the administrator can schedule how often an ACO occurs for each individual provider (weekly, monthly, bi-monthly, and so on), that administrator can predict how often profiles will be affected and reduce how frequently providers 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 provider, a local cache copy of the profile is on the *Dragon Medical Client* workstation. Then, on a typical day, when a provider opens his or her user profile on the *Dragon Medical Client* in a *DM360 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 provider 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 1.3 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 or Call Forwarder 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.

Summary of Test Environments

Below are additional details on test environments that produced the data provided in this chapter.

Test Environment for Estimating Dragon Response Times: Hardware and Software Configurations for Servers, Clients, and Simulator

Hardware Item	NMS Server	Dragon Medical Client	Simulator
Processors	Two 3.349 GHz Pentium 4 Processors	2.79 GHz Pentium 4	2.66 GHz Pentium 4
Memory	4 GB	2 GB	2 GB
Operating System	Windows 2003	Windows XP Pro SP3*	Windows XP Pro SP3*
Other Software	SQL Server 2008	*Dragon Medical Client* Version 10.1	
Network Bandwidth	100 Mbps		

To simulate multiple provider activities on *Dragon Medical Clients* that involve *NMS Server*, the test used a home grown simulator and related scripts. The test generated these provider activities:

- 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 Medical 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	*Dragon Medical Client* vSync Version 10.1	*Dragon Medical Client* 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)
Operating System	Windows Server 2003 R2 Standard Edition SP2 32-bit	Windows XP SP3*
Other Information		40 Virtual machines running *Dragon Medical Client* software.

* The 2.x versions of DM360 Network Edition will be the last versions of the product to support Windows XP. DM360 Network Edition 2.1 continues to support Windows XP.

Index

A

accent in speech
- effect on how often to run ACO ... 19

ACO process ... 18
- defined ... 18
- how it differs from LMO ... 18
- how often to run ... 19
- maintenance ... 19
- minimum recommended frequency ... 19
- recommended frequency ... 19

adaptation
- defined ... 18

B

basic components of DM360 Network Edition ... 8

Bluetooth wireless microphone support ... 60

C

calculation
- number of minutes for ACO ... 22
- number of nodes required ... 21
- percentage of providers dictating ... 23
- space for local cache user profiles ... 16
- space for master user profiles ... 14

calculation example
- space for master user profiles ... 14

calculation vs. rule of thumb
- number of nodes for total providers ... 23

D

databases
- storage space for NMS Server ... 25
- storage space for Prof Opt Server ... 25

dictation volume
- effect on how often to run ACO ... 19
- effect on how often to run LMO ... 20
- how often to run LMO based on ... 20

DM360 Network Edition
- network components ... 8
- overview ... 8

Dragon logs
- how often purged ... 26

duplex settings on NICs ... 27

E

errors in recognition
- effect on how often to run ACO ... 19

F

features of DM360 Network Edition ... 7

H

hardware/software requirements
- master user profiles storage ... 44

hardware/software requirements summary
- network configurations ... 47

I

improved recognition
- how ACO results affect ... 21

IP switch
- load balancing on NMS Servers ... 27

K

knowledge
- required to plan network ... 3

L

large configuration
- hardware and software details ... 43
- network layout for 1,000 providers ... 37

link speed settings on NICs ... 27

LMO process ... 18
- defined ... 18
- how it differs from ACO ... 18
- how often to run ... 20

load balancing
- multiple NMS Servers ... 27

local cache user profiles
- how ACO results integrate with ... 21
- space requirements calculation ... 16

log files
- storage space requirements ... 26

logs
how often purged 26

M

master user profiles
space requirements calcualtion 12
master user profiles machine
hardware/software requirements 44
medium configuration
hardware and software details 41

N

network configurations
Dragon Medical 360 | Network Edition 33, 52
hw/sw requirements summary 47
options available 33, 52
summary table of servers 50
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 18
providers
how often to run LMO for experienced 19
how often to run LMO for new 20
Speech Nodes required by number of 23

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 40
Speech Node service logs
how often purged 26
Speech Nodes
calculating number required 21
number based on total providers 23
role in ACO and LMO processes 18
speech patterns
processes for learning 18
speed of network equipment 27
storage space
requirements for log files 26
requirements for NMS database 25
requirements for Prof Opt database 25
Storage space
Server installations 26
storage space requirements
calculation for local cache user profiles 16
calculation for master user profiles 12
summary table
Network configurations 50

T

tuning acoustic and language models 18

U

Update Server 20
user profiles
number of nodes for optimization 21
optimization 18
space requirements calculation 12

V

vocabularies
effect on how often to run LMO 20

W

words
effect on how often to run LMO..... 20