

Level: Intermediate
Works with: Notes/Domino
Updated: 04-Nov-2002

by
Julie
Kadashevich

Getting an agent to run means much more than just writing it. As we discussed in a previous article, "[Controlling the agents in your system](#)," there are various rights and restrictions that determine whether or not you can run an agent. Once you've got the appropriate rights, you're all set if you run your agent manually (for example, from the Actions menu or from a button on a form). However, there are more factors to consider if you've designed an event-triggered or scheduled agent because it's up to the Agent Manager to run your agent.

There are a number of ways in which agents can be invoked both on the server and the client:

- The Agent Manager server task controls *scheduled* agents. These include agents that run on a schedule specified by the user or are scheduled to run when an event occurs (such as when new mail arrives or when a new document is created).
- The HTTP task runs agents invoked via a browser.
- The router runs agents executed as part of mail delivery ("before mail delivery" agent trigger).
- The Domino server runs agents invoked via the RunOnServer command, issued from agents invoked on the client. Agents can also be invoked on the server through server console commands.
- The Notes client runs agents invoked from the client. The Client TaskLoader runs scheduled agents on the client, using an embedded version of the Agent Manager.

HTTP agents, router agents, agents invoked via RunOnServer, and nonscheduled client agents run immediately. In this article we will concentrate on scheduled and event-triggered agents run by the Agent Manager. By using the Agent Manager parameters that control the scheduling of agents, you can control server resources. The downside is that to utilize resources more efficiently for all server tasks, the agents do not run immediately.

This article will help you determine when the Agent Manager will run your agent and discusses possible ways to minimize any delays. We'll focus on event-triggered agents, or agents that are triggered by such things as new mail arriving in a database or by a document being created or modified. The execution of these agents depends on how quickly the Agent Manager queues the agents, when the Agent Manager schedules them to run, how heavy the load is on the Agent Manager, and how many concurrent agents are running. We'll finish up by looking at an example of the minimum delays possible.

One thing to remember is that although we'll discuss ways to modify settings for the Agent Manager, the default settings are usually recommended. The best values for your situation will depend on your particular server and its specific load.

When will the Agent Manager queue my agent?

The Agent Manager periodically checks to see if it has any new agents that it needs to schedule. The following two

environment variables control how quickly an agent gets into the schedule queue:

- AMgr_SchedulingInterval specifies a delay (in minutes) between running of the Agent Manager's scheduler. Valid values: one minute to 60 minutes. The default value is one minute.
- AMgr_UntriggeredMailInterval specifies a delay (in minutes) between running of the Agent Manager's check for untriggered mail. (Mail is "untriggered" when it's added to a database through replication, which does not cause a new mail trigger to fire. This parameter helps guarantee that the replicated mail will be processed.) Valid values: one minute to 1440 minutes (the number of minutes in a day). The default value is 60 minutes.

The Agent Manager automatically uses the default values for these two variables. To override the default values, you can add these variables to the server's Notes.ini file. But, in general, these two variables rarely require tuning.

When will the Agent Manager execute my agent?

When you create or modify an event-triggered agent, the Agent Manager schedules it to execute right away so the agent can process any new documents up to that point. The subsequent executions are controlled by the following two sets of environment variables that can add time into the agent's schedule—affecting how quickly the Agent Manager will execute the agent again.

The first two environment variables specify a delay between when an event occurs (a document is updated or new mail is delivered) and the time at which the agent is scheduled to be executed. This delay exists so a rapid series of events will only result in one execution of a given agent. Then, the Agent Manager as well as the overall server can perform more efficiently.

- AMgr_DocUpdateEventDelays specifies a delay of execution (in minutes) that the Agent Manager includes in the schedule for a document update-triggered agent after a document update event. The default value is five minutes.
- AMgr_NewMailEventDelay specifies a delay of execution (in minutes) that the Agent Manager includes in the schedule for a new mail-triggered agent after a new mail message is delivered. The default value is one minute.

The next two environment variables specify a delay between the execution of the same event-triggered agent (triggered by a document update or by new mail). This delay exists so the user has control over the interval between the given agents. If the agent is triggered and its corresponding minimum interval hasn't passed, the Agent Manager schedules the task to run at its last run time plus the minimum interval.

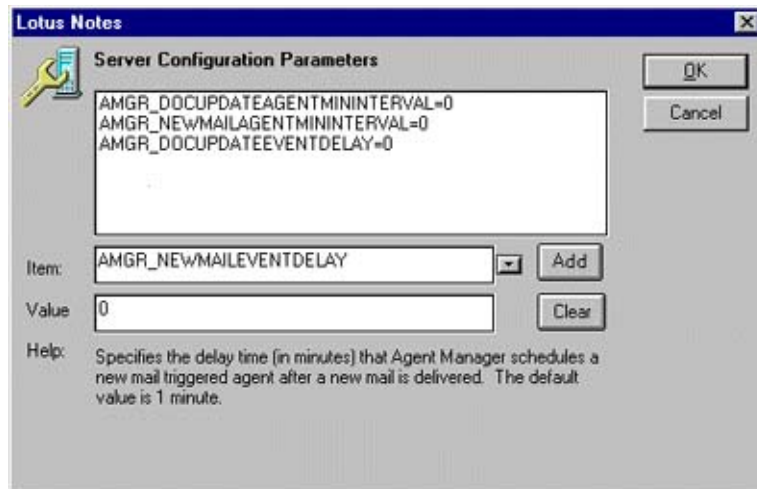
- AMgr_DocUpdateAgentMinInterval specifies the minimum elapsed time (in minutes) between execution of the same document update-triggered agent. The default value is 30 minutes.
- AMgr_NewMailAgentMinInterval specifies the minimum elapsed time (in minutes) between execution of the same new mail-triggered agent. The default value is zero.

If you find that you need to run agents more often than what these default values specify, you can tune the Agent Manager by reducing the values. If you would like to reduce the system resources used by the Agent Manager, you can increase the values.

Setting these variables in the Server Configuration document

Although you can add these variables directly to the server's Notes.ini file, an easier way is to simply edit the Server Configuration document in the Domino Directory. It's easier because you can use a Server Configuration document to specify Notes.ini settings for a single server, a group of servers, or for all the servers in a domain. These settings are then written to the Notes.ini file.

1. Open the Domino Directory, and go to the Server - Configurations view. If a Server Configuration document already exists for your server, select it and click Edit Configuration. Otherwise, click Add Configuration.
2. Select the Basics tab. In the Group or Server name field, enter the server name. (You can also enter the name for a group of servers or an asterisk * to modify all servers in the domain.) Then open the NOTES.INI Settings tab, and click Set/Modify Parameters.
3. Next to the Item field, click the arrow to display a list of standard parameters, and choose one. (Notice that a brief help description appears for the parameter.)
4. Enter a value for the setting.
5. Click Add to confirm the value, and then choose another parameter. Click OK when finished. Then, click Save and Close.



In order for these changes to take effect, the Agent Manager cache needs to be refreshed. You can either wait for the next automatic refresh (as specified in the Server Tasks - Agent Manager tab of the Server document, under Basics - Refresh agent cache), which by default is at 12 a.m. Or, you can restart the Agent Manager by entering the following commands at the server console:

```
tell amgr quit  
load amgr
```

You can also restart the Agent Manager by restarting the server.

How does the Agent Manager's load affect agent execution delays?

Execution of all agents (event-triggered agents as well as scheduled agents) also depends on how busy the Agent Manager is and how you've allocated system resources for the Agent Manager. (If you are running Domino 4.6, see the [Domino 4.6 specific information](#) sidebar for details about the "Max % busy before delay" parameters.)

If your system is attempting to schedule agents at a rate faster than the Agent Manager can execute them, the Agent Manager notices that its internal queues are growing and generates the following message:

AMgr: Agent scheduling is paused

This message means that the Agent Manager will not schedule any new agents until it processes some of the agents that were already scheduled. The new agents will not be lost, but they may run with a delay.

These warning messages could be an indication that the number of allowed concurrent agents is too low for the number of agents that your server is trying to process. To improve the performance of the Agent Manager, you can increase the number of allowed concurrent agents

How many concurrent agents are running?

One of the ways to relieve a heavily loaded Agent Manager is to allow concurrent agent execution. To do this, you can modify the Max concurrent agents field in the Agent Manager section of the Server document. Increasing the number of concurrent agents allows more than one agent to run at the same time. The valid values for the Max concurrent agents field are one to ten. The default values are one for daytime and two for nighttime.

Each concurrent agent is executed by a separate Agent Executive. (The terms *concurrent agents* and *multiple Agent Executives* refer to the same concept.) To see a snapshot of the Agent Manager status, including how many Agent Executives are currently running, you can enter the following command at the server console:

```
tell amgr status
```

Running each agent in its own process is a fault-tolerant paradigm because if one of the agents fails, it does not affect the other agents.

Increasing the value of the Max concurrent agents field improves the responsiveness of the Agent Manager. Because all resources in the system are shared, giving more resources to the Agent Manager may affect the performance of the rest of the server.

Tip: When a new agent needs to be scheduled, the Agent Manager checks to see whether or not any Agent Executives are already executing agents in the applicable database. If so, the new agent is queued to the same Agent Executive for efficiency. Otherwise, the next available Agent Executive starts executing the agent. As a result, if you have numerous agents residing in one database, they do not run concurrently unless you place them into multiple databases.

What are the minimum delays?

So far, you've learned about the Agent Manager delays that you can control by modifying environment variables and Agent Manager fields in the Server document. In addition, the Agent Manager has several small delays built into the system. These small delays ensure that the Agent Manager does not consume all the CPU cycles if all user-defined parameters are set to zero.

To answer the question "What are the minimum delays?" I timed the delays between an event and the agent execution on a Pentium Pro200 system with a low load. I used the following settings in my Notes.ini file:

- AgentManagerSchedulingInterval = 1
- AMgr_DocUpdateAgentMinInterval = 0
- AMgr_NewMailAgentMinInterval = 0
- AMgr_DocUpdateEventDelay = 0
- AMgr_NewMailEventDelay = 0

For agents that were triggered by a document update, the average time between document creation and agent execution was approximately two minutes.

For agents that were triggered by new mail, the average time between new mail arrival and agent execution was approximately two minutes.

These numbers varied somewhat from execution to execution on my PC because the Agent Manager is affected by other tasks that the server has to perform.

Conclusion

In this article, we examined the factors that affect how quickly the Agent Manager executes your event-triggered agents. Because the tasks performed by each server are different, it may require some tuning on the administrator's part to get all the settings into the optimal configuration for a particular system. I hope this article gave you a better understanding of how you can achieve the best results.

ABOUT THE AUTHOR

Julie Kadashevich came to Iris in March of 1997 after working at FTP Software on Java and C++ mobile agent technology. This is her sixth article for *LDD Today*. She focuses on the Agent Manager as well as Java. Previously, she worked in the area of applied Artificial Intelligence at Wang Labs and received five patents in the field. She has Bachelor's and Master's degrees from Boston University, both in computer science.



Domino 4.6 specific information

In Domino 4.6, the "Max % busy before delay" parameters in the Agent Manager section of the Server document specify the approximate percentage of time the Agent Manager can use to run agents. If the Agent Manager exceeds this percentage, it delays agent executions until the percentage falls below the maximum. The valid values for the "Max % busy before delay" field are 0 percent to 90 percent. The default values are 50 percent for daytime, and 70 percent for nighttime.

If the Agent Manager exceeds the maximum percentage of time allowed, it delays the agent execution and displays the following message on the server console:

AMgr: Percentage of time agents allowed to run exceeded, delaying additional agent executions

This warning message could be an indication that the "Max % busy" field on your server is set too low for the number of agents that your server is trying to process. To improve the performance of the Agent Manager, you can increase the value of the "Max % busy" field. Since all resources in the system are shared, giving more resources to the Agent Manager may affect the performance of the rest of the server.

Note that "Max % busy before delay" does not apply to Domino 5 and 6.