

Contents

- 1 Scenario Setup
- 2 Problem Statement
- 3 Isolation
 - ◆ 3.1 Email Server / Retriever Process
 - ◆ 3.2 Rules Process
 - ◆ 3.3 External Agent Assignment Service (EAAS) Process
- 4 Tasks in Queue
- 5 Check MRD Configuration
- 6 Agent Concurrent Task Limit
- 7 Resolution

Scenario Setup

UCCE

1. Set MAX CALLS IN QUEUE for MRD_EIM to 2 (from documented requirement=5000)

Administration Console

1. Set integrated agent's concurrent_task_limit for the EIM queue to 2

System Console

1. Stop the EAAS process
2. Start the EAAS process and instance

Agent PC

1. Send 5 emails to support@eim.lab
2. Log in as integrated agent "goofy" and complete all emails in queue. (Can also use a second agent "mmouse" if desired to speed up the activity completion)

Problem Statement

It is taking an unusual amount of time for incoming emails to reach agents. A handful of emails sent to an empty queue are taking several minutes to fully process. We first need to rule out some components if possible to be able to isolate from where the problem is coming.


Isolation

Email Server / Retriever Process

How quickly are the emails being pulled off of the email server? Recall that the Retriever process polls the email server (by default) every 30 seconds, pulling 10 emails with each pass.

1. We can verify the origination time on the mail server. In Exchange, this is the Message Tracking Center.

Delays_in_Email_Delivery_to_Agents



2. We can then verify two things in the Retriever logs at INFO level: (1) Retrieval time and (2) Time between "Waiting" and "Running" messages.

```
2011-06-22 15:08:31.216 GMT-0400 <@> INFO <@> [47:RxInstance id : 999] <@> ProcessId:1332 <@> PID:
2011-06-22 15:09:01.341 GMT-0400 <@> INFO <@> [47:RxInstance id : 999] <@> ProcessId:1332 <@> PID:
2011-06-22 15:09:01.341 GMT-0400 <@> INFO <@> [47:RxInstance id : 999] <@> ProcessId:1332 <@> PID:
2011-06-22 15:09:01.373 GMT-0400 <@> INFO <@> [47:RxInstance id : 999] <@> ProcessId:1332 <@> PID:
2011-06-22 15:09:01.373 GMT-0400 <@> INFO <@> [47:RxInstance id : 999] <@> ProcessId:1332 <@> PID:
```

If the Retriever is successfully polling the mail server every 30 seconds and pulling emails off of it, we can rule those two components out.

Rules Process

Once the email is retrieved and becomes an activity, it must go through the inbound workflow via the Rules Process, which also runs every 30 seconds. With rules-process logs at INFO-level tracing, you can track the nodes an activity goes through to see if any delays are being introduced. From the below log snippets we can see that our three emails are being processed within two runs of the rules-process from when they were retrieved. Our delays in reaching the agent are much more than ~45 seconds, so we can rule out the Rules Process and our inbound workflow.

```
2011-06-22 15:09:36.514 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.530 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.561 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.639 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.717 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.780 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.842 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
2011-06-22 15:09:36.873 GMT-0400 <@> INFO <@> [68:BPRulesInstance id : 996] <@> ProcessId:2192 <@>
```

External Agent Assignment Service (EAAS) Process

Once the activity is processed through the workflow, EAAS will pick it up and request routing instructions from UCCE. The EAAS logs can be reviewed for NEW_TASK messages to see if there is some delay being

Delays_in_Email_Delivery_to_Agents

introduced here.

```
2011-06-22 15:13:00.643 GMT-0400 <@> INFO <@> [43:pool-2-thread-1] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:00.659 GMT-0400 <@> INFO <@> [44:pool-2-thread-2] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:21.018 GMT-0400 <@> INFO <@> [68:pool-2-thread-3] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:21.971 GMT-0400 <@> INFO <@> [73:pool-2-thread-4] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:34.972 GMT-0400 <@> INFO <@> [75:pool-2-thread-5] <@> ProcessId:2612 <@> PID:1 <@>
```

NEW_TASKS appear to be sending slowly for consecutive activity IDs. They also seem to be going in pairs. What about the corresponding **DO_THIS_WITH_TASK** messages?

```
2011-06-22 15:13:21.284 GMT-0400 <@> INFO <@> [67:pool-4-thread-1] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:21.971 GMT-0400 <@> INFO <@> [72:pool-4-thread-2] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:34.972 GMT-0400 <@> INFO <@> [74:pool-4-thread-3] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:43.956 GMT-0400 <@> INFO <@> [81:pool-4-thread-4] <@> ProcessId:2612 <@> PID:1 <@>
2011-06-22 15:13:51.956 GMT-0400 <@> INFO <@> [88:pool-4-thread-5] <@> ProcessId:2612 <@> PID:1 <@>
```

Some pretty significant delays between **NEW_TASK** and **DO_THIS_WITH_TASK**: 30 seconds, with only 5 emails in queue!

Tasks in Queue

With INFO-level EAAS logs, you can see how EIM keeps track of the number of tasks in queue relative to the maximum number of tasks allowed. Look in your logs for the below "incrementCount" and "decrementCount" messages.

Note that the value "concurrent_task_limit" in the incrementCount/decrementCount messages reflects the current number of activities queued to an MRD, and not any sort of limit. Cosmetic defect CSCtq55038 has been filed to have this part of the log message renamed appropriately (i.e. "tasks_queued").

```
2011-06-22 15:12:34.111 GMT-0400 <@> INFO <@> [29:Thread-6] <@> ProcessId:2612 <@> PID:1 <@> UID:1
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 1 : Type = java.lang.Integer
<@>
2011-06-22 15:12:34.127 GMT-0400 <@> INFO <@> [29:Thread-6] <@> ProcessId:2612 <@> PID:1 <@> UID:1
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 2 : Type = java.lang.Integer
<@>
2011-06-22 15:13:20.987 GMT-0400 <@> INFO <@> [67:pool-4-thread-1] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 1 : Type = java.lang.Integer
<@>
2011-06-22 15:13:20.987 GMT-0400 <@> INFO <@> [67:pool-4-thread-1] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 2 : Type = java.lang.Integer
<@>
2011-06-22 15:13:21.956 GMT-0400 <@> INFO <@> [72:pool-4-thread-2] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 1 : Type = java.lang.Integer
<@>
2011-06-22 15:13:21.956 GMT-0400 <@> INFO <@> [72:pool-4-thread-2] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 2 : Type = java.lang.Integer
<@>
2011-06-22 15:13:34.956 GMT-0400 <@> INFO <@> [74:pool-4-thread-3] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 1 : Type = java.lang.Integer
<@>
2011-06-22 15:13:34.956 GMT-0400 <@> INFO <@> [74:pool-4-thread-3] <@> ProcessId:2612 <@> PID:1 <@>
```

Delays_in_Email_Delivery_to_Agents

```
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 2 : Type = java.lang.Integer
<@>
2011-06-22 15:13:43.956 GMT-0400 <@> INFO <@> [81:pool-4-thread-4] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 1 : Type = java.lang.Integer
<@>
2011-06-22 15:13:51.956 GMT-0400 <@> INFO <@> [88:pool-4-thread-5] <@> ProcessId:2612 <@> PID:1 <@>
max_task_limit = 2 : Type = java.lang.Integer
concurrent_task_limit = 0 : Type = java.lang.Integer
<@>
```

Check MRD Configuration

From the EAAS logs we can see that the maximum number of tasks allowed to be queued to MRDID 5003 is 2. This value, "max_task_limit", corresponds to the setting "Max Calls in Queue" configured at the MRD level from the UCCE Configuration Manager. Per the Deployments Guide, this value is incorrect and is recommended to be set to 5000.

The screenshot shows the configuration for MRD_EIM. The Name is MRD_EIM, Media routing domain ID is 5001, and Media class is CIM_EIM. The Task section includes Life (5000 seconds), Start timeout (30 seconds), and Max duration (28800 seconds), all with 'Override Media Class Default' checkboxes. The Calls in Queue section shows Max (2), Max per call type, and Max time in queue. Service level threshold is 30, Service level type is Ignore Abandoned Calls, Interruptible is checked, and Description is MRD for inbound eMail.


Field	Value	Unit	Override Media Class Default
Name	MRD_EIM		
Media routing domain ID	5001		
Media class	CIM_EIM		
Task Life	5000	seconds	<input type="checkbox"/>
Task Start timeout	30	seconds	<input type="checkbox"/>
Task Max duration	28800	seconds	<input type="checkbox"/>
Calls in Queue Max	2		
Calls in Queue Max per call type			
Calls in Queue Max time in queue		seconds	
Service level threshold	30		
Service level type	Ignore Abandoned Calls		
Interruptible	<input checked="" type="checkbox"/>		
Description	MRD for inbound eMail		

With the above information, we can see that UCCE will not allow more than 2 tasks to be queued at a time. Therefore, while those 2 tasks are queued, EIM will not send any additional NEW_TASK messages for routing.

EAAS Process and Instance must be cycled for changes to take effect.

Agent Concurrent Task Limit

A low value for MRD Max Calls in Queue will also impact the concurrent task limit that can be set for an agent. In this scenario, our agent's concurrent task limit was set to 2 as well. If we try to set it to 3, we receive an error:



Resolution

The above scenario demonstrated the difference between `max_task_limit` and `concurrent_task_limit`. Now consider a production environment with the same "Max Calls in Queue" that receives more than 5 emails in a few minutes. If 10 emails are retrieved, only **two** will be be queued at a time. Agents will still be able to reach their max tasks, but it will take much longer to do so.

An incorrectly configured "Max Calls in Queue" can become a serious problem as incoming email and queue volume increases, or in the event of EAAS failover when all queued activities must be re-submitted for routing.