

# Schedule Definition

Custom Software Created By



Documentation By  
Kevin King  
Precision Solutions, Inc.  
303/651-7050

Last Updated  
3/9/2007 9:14:00 PM

## Introduction

The Schedule Definition tool is a custom tool for SB+ developers that wraps around the existing SB+ Job Schedule Definition tool and provides recurrent scheduling of background jobs. Analogous to the Unix `cron` verb – but with a user-friendly interface! – the Schedule Definition tool allows SB+ processes to be automatically restarted on very flexible schedules.

While SB+'s Job Schedule Definition tool enables an SB+ process to be run on a background (phantom) session, a process can be invoked only once per background session. There is no functionality incorporated into the standard SB+ tool that allows a job to be rescheduled for regular execution. **(While some have tried to create a job schedule that invokes itself, this works only for a few iterations before the recursive session runs out of memory and begins behaving improperly.)**

The Schedule Definition tool fills this gap by providing a "scheduler for job schedules", that is, a tool for running job schedule definitions (and other processes) at regularly scheduled intervals. By setting up the processes you want to run on certain intervals, there's no need to monitor and manually restart your SB+ processes; let the computer do the work for you!

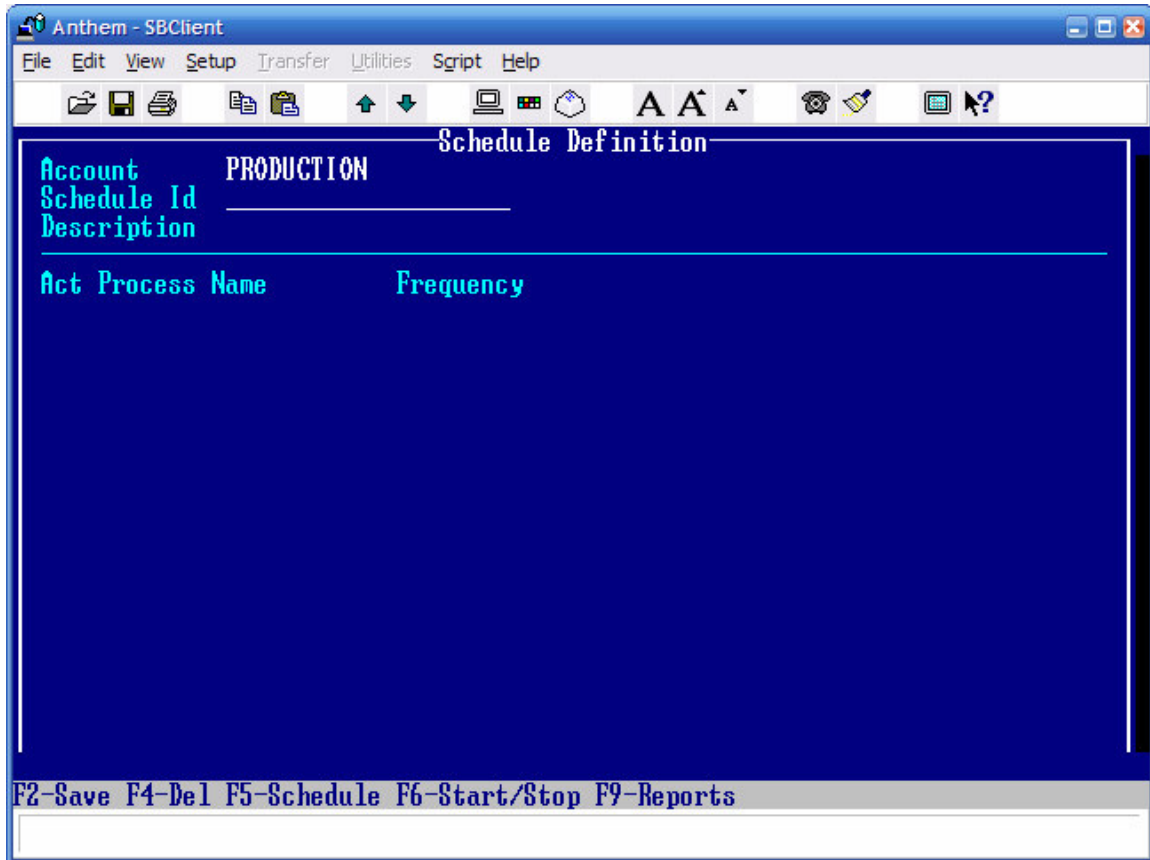
## Starting/Stopping the Schedule Definition Tool

Schedule Definition is a one-stop-shopping tool for creating, editing, modifying, deleting, starting, and stopping all of your background processes. Enter `/PJSD` at any menu or input prompt to start this tool. Once started, the screen shown in the following section will be displayed.

To exit the tool, simply press `<esc>` when asked to enter a Schedule Id.

## The Schedule Definition Screen

The /PJSD screen is where you define the processes you want to run and the intervals for each. When the tool is started, the screen appears as follows:



At the Schedule Id prompt, enter the name of a schedule that you want to create, edit, or delete. To see a list of schedules currently defined, press F3 at this prompt. Once a schedule Id has been entered, the details of that schedule will be displayed.

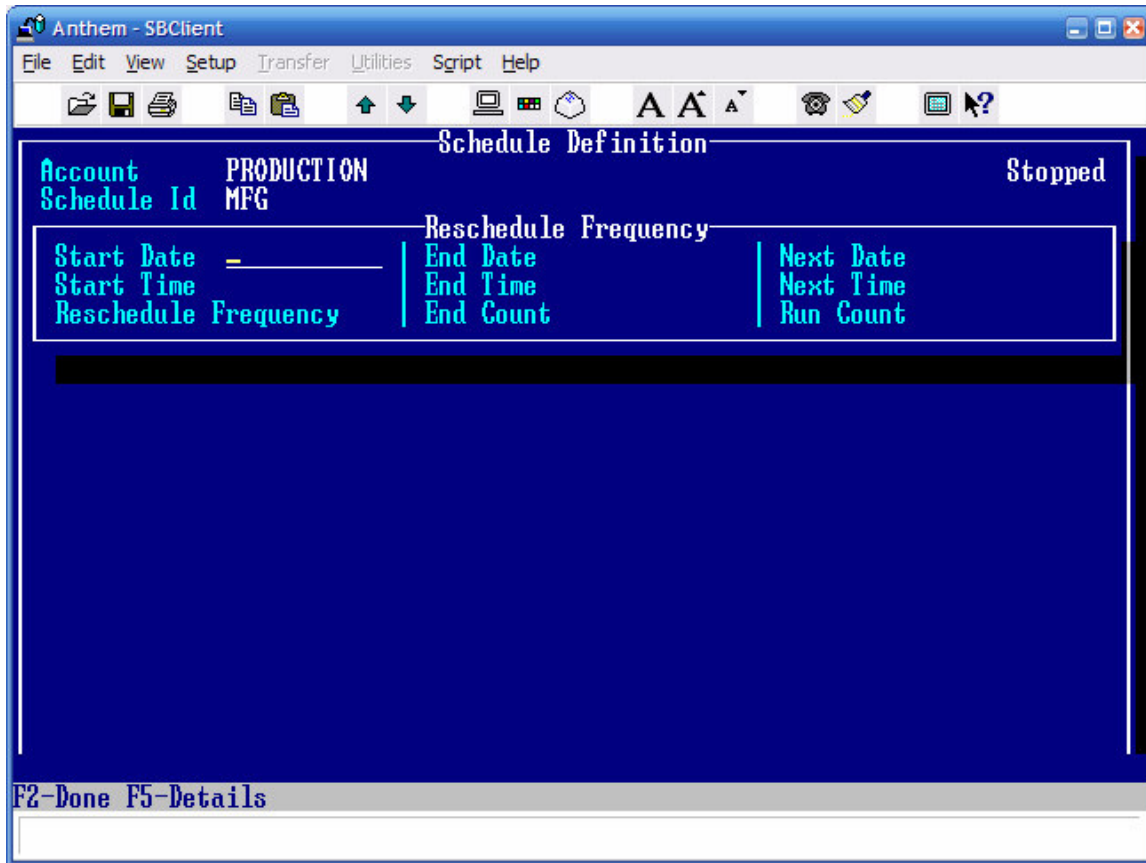
At the Description prompt, enter the description for this schedule. The description entered here appears only when you use F3 on the previous prompt, and is used for finding schedules at some later point in time.

The Act, Process Name, and Frequency prompts allow you to define the processes you want to run and the interval for each.

Act (active) is either Y or N. If a process is active, it will be run at regular intervals (described by the text shown in the Frequency prompt). If a process is inactive, it will not be started by the scheduler. Using this flag you can dynamically enable and disable processes while the scheduler is running.

At the Process Name prompt, enter the name of a process to be run on a regularly scheduled interval. Any process available in the current system can be entered here, not just regular SB+ job schedule processes. Note, however, that if the process to be run requires any user input you might create a regular job schedule definition to wrap around that process to provide the required inputs. Then, run the job schedule process via this tool.

The Frequency prompt displays the currently defined interval for this job. When you first add a process to the schedule, there is no frequency defined. Therefore, another screen pops up to gather this information:



At the Start Date prompt, enter the first date you want this process to be run. Leave this (and the Start Time) blank if you want the process to be run immediately when the scheduler is started.

At the Start Time prompt, enter the first time you want this process to be run. (If the Start Date is not entered, you will not be allowed to enter a start time.)

The Reschedule Frequency allows you to define how regularly the process should be run. Valid options here include:

- Once Only
- Seconds
- Minutes
- Hourly
- Daily
- Weekly
- Monthly
- Yearly

You may press F3 at this prompt to see this list of choices. Based on your choice at this prompt, one of several additional screens will appear asking for additional information. You will find more information about these choices and the options available for each later in this document.

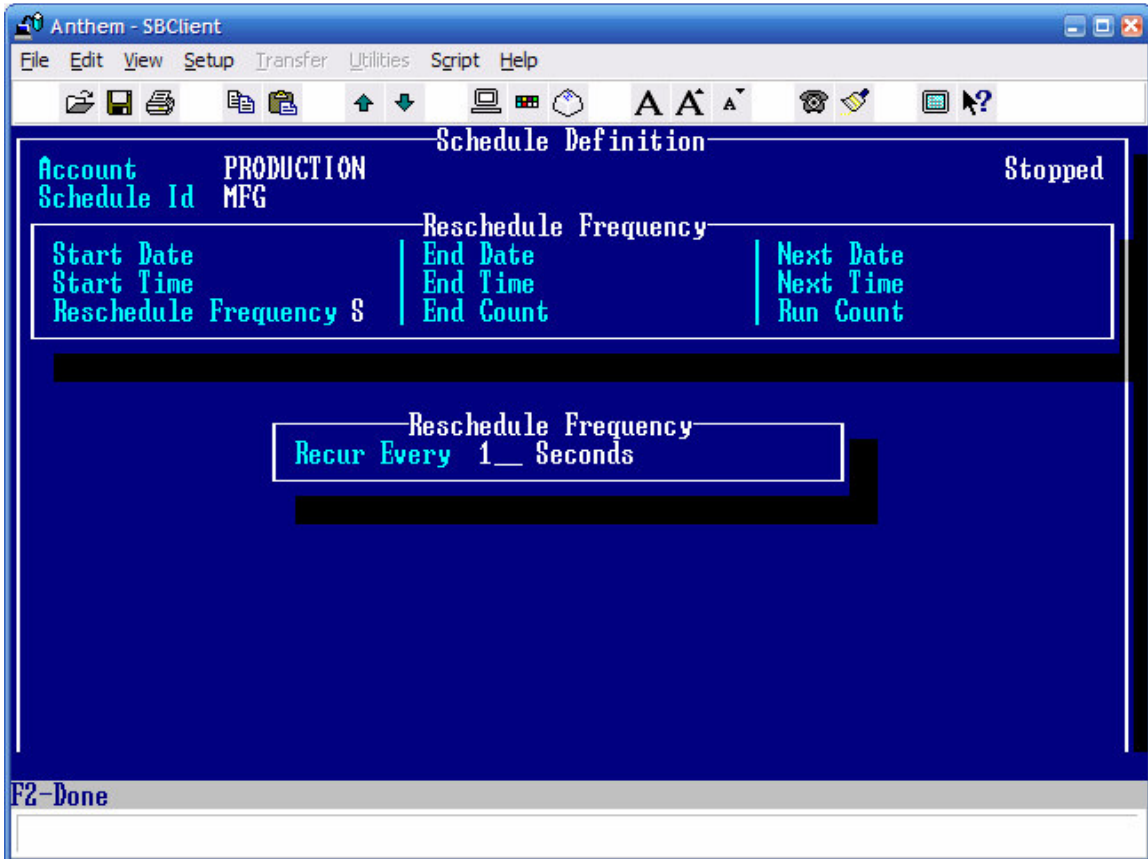
The End Date, End Time, and End Count prompts define when this process will stop running. If you want the process to run up to a specific date and time, enter that date and time here. If you want the process to run only a certain number of times, enter that value into the End Count prompt.

The Next Date, Next Time, and Run Count prompts are for your information only. The scheduler will fill in these values when it runs the process and calculates the next date and time the process will be run. Note that the process will run only when the schedule is running. For more information see "Starting/Stopping the Schedule" later in this document.

Once a process has been run once, you will not be allowed to change the start date, start time, or reschedule frequency for that line. You can, however, press F5-Details to view the details of that particular scheduled item. (If you need to change the reschedule frequency for a given process, you can deactivate the original process on the schedule and add the process again on a new line in with a new frequency code and associated details.)

## Scheduling a Second-by-Second Process

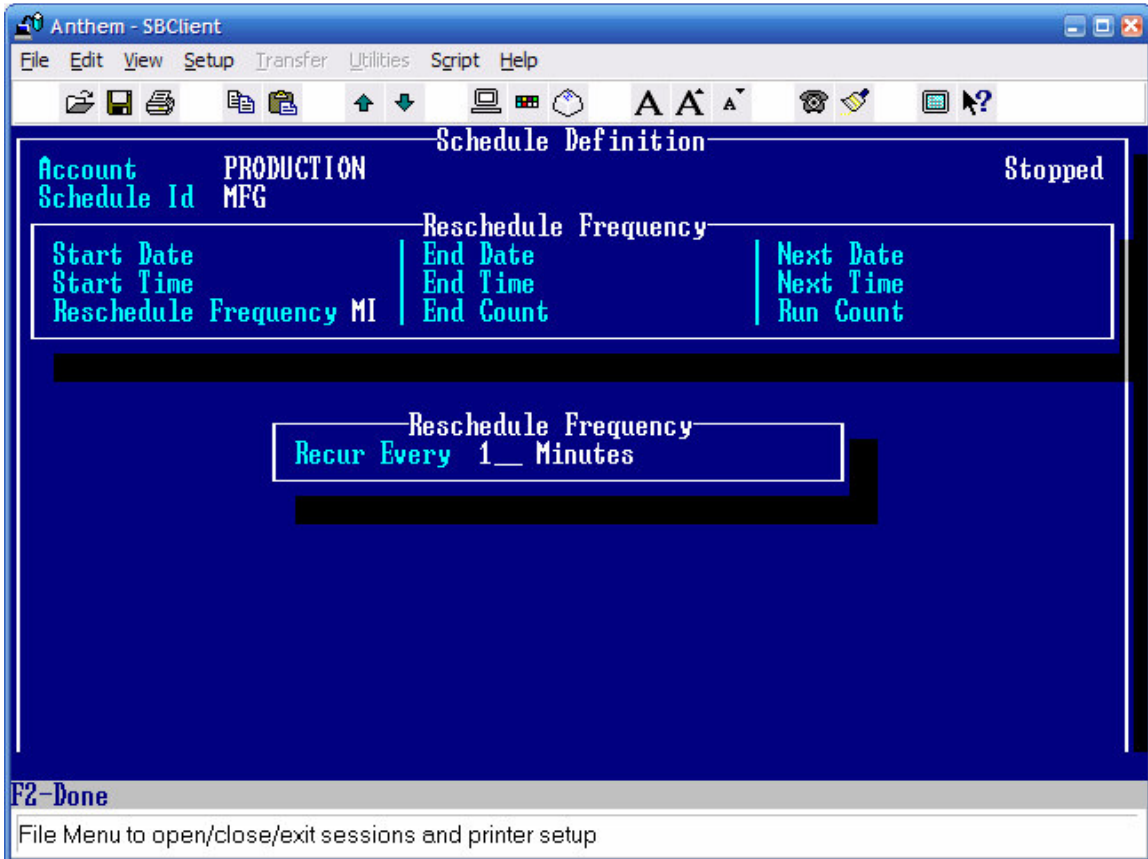
For a given process, when prompted for Reschedule Frequency, enter "S". The following screen will appear.



At this prompt, enter the number of seconds between runs. Note that there is generally about a 5-10 second delay in starting any job, so anything less than 10 will result in the job running constantly.

## Scheduling a Minute-by-Minute Process

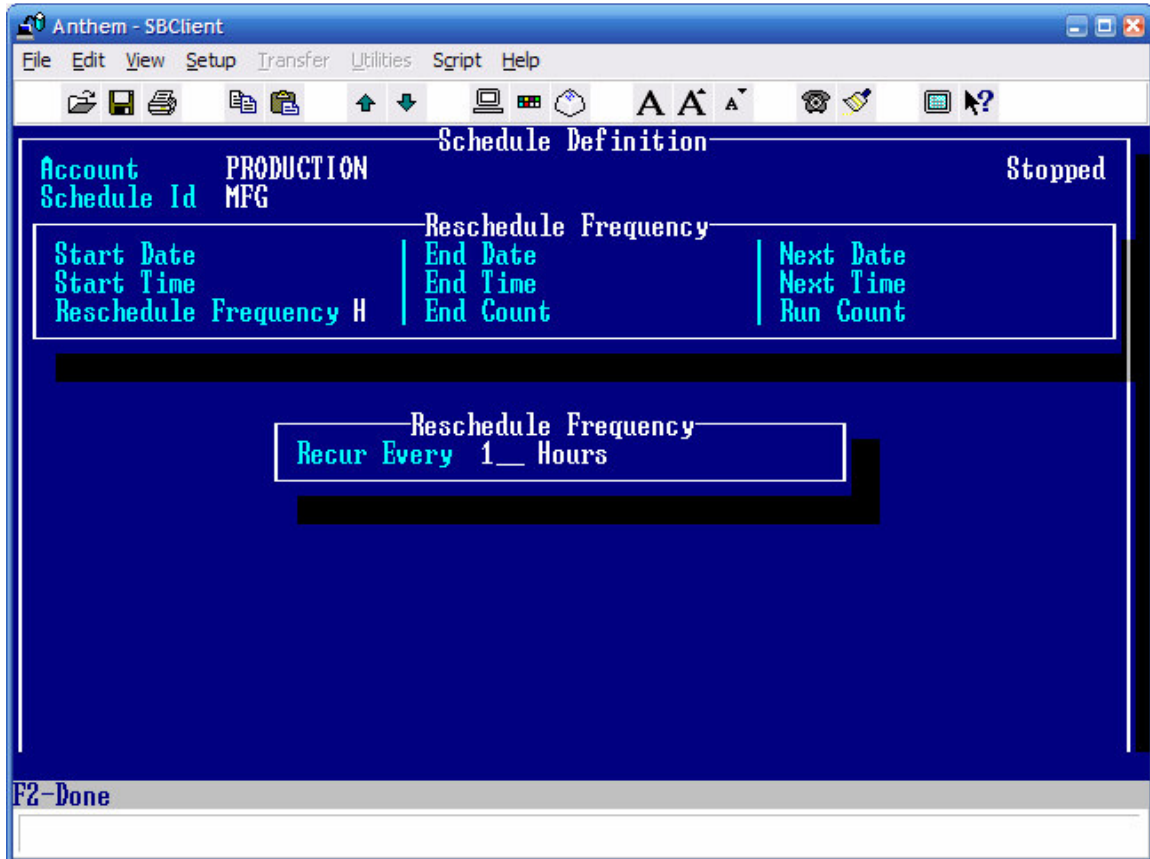
For a given process, when prompted for Reschedule Frequency, enter "MI". The following screen will appear.



At this prompt, enter the number of minutes between runs. To run once every minute, enter 1. To run twice an hour, enter 30. For four times an hour enter 15.

## Scheduling an Hourly Process

For a given process, when prompted for Reschedule Frequency, enter "H". The following screen will appear.

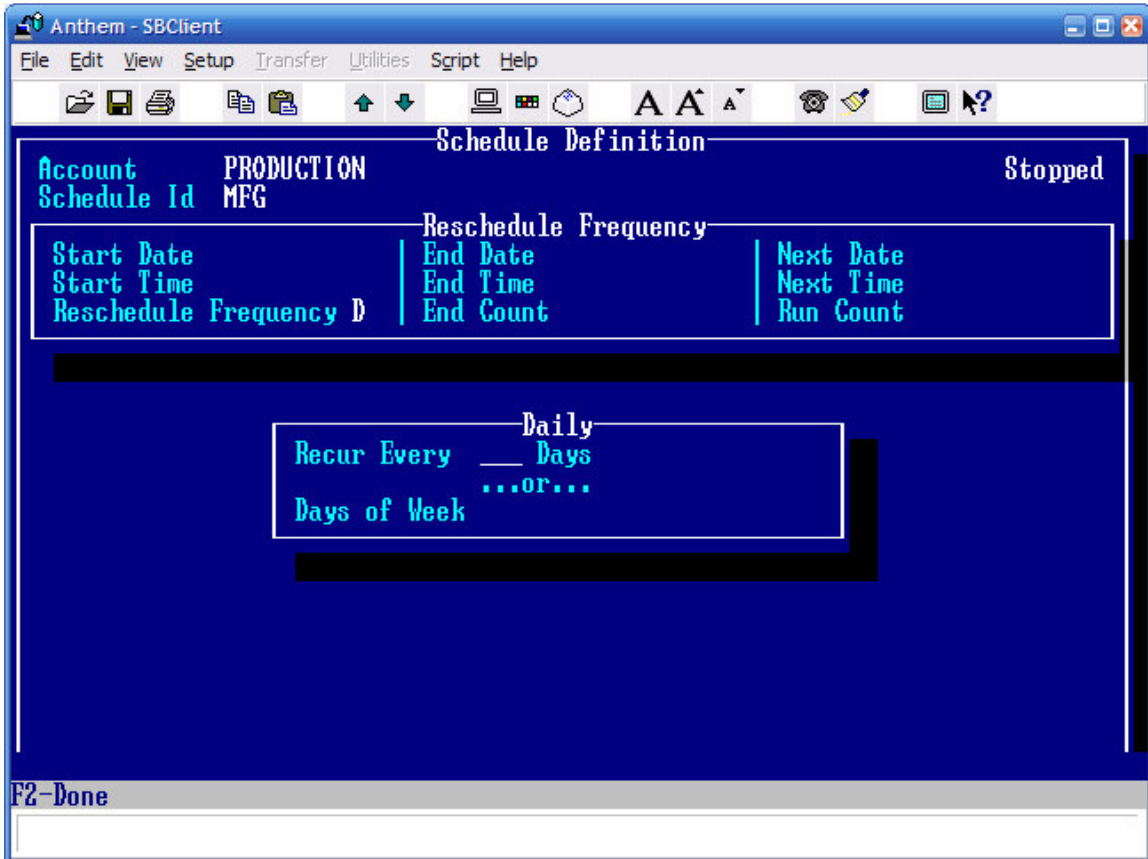


At this prompt, enter the number of hours between runs. If you want the process to run every hour, enter 1. For every other hour, enter 2. To run the job once in the morning and once in the evening, enter 12. Any number of hours can be entered here, though at some point one of the other options (daily, weekly, etc.) might be a better choice for the interval.

Press F2 to save the information entered into this screen.

## Scheduling a Daily Process

For a given process, when prompted for Reschedule Frequency, enter "D". The following screen will appear.



In this screen, you can enter either a number of days between runs, or the days of the week you want the process to be run.

At the first prompt, enter a number that defines the number of days between runs. To run every day, enter 1 here. To run every other day, enter 2 here. To run once a week, you may enter 7 here (however, at this point the weekly interval may be a better choice.)

If you leave the first prompt blank, you can then enter the days of the week for running this process. At this prompt you may press F3 to select the days, or enter the codes for the days you want to use, using the following table:

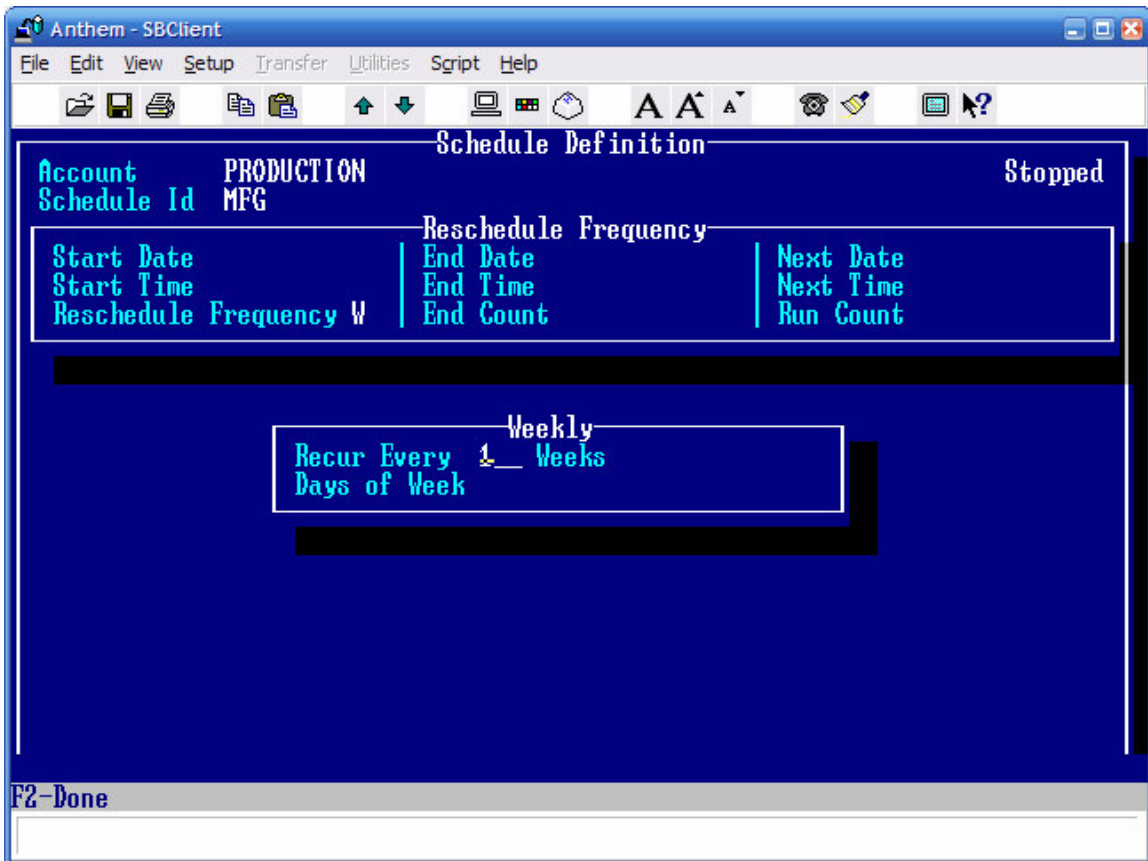
For this day...	Use this code.
Monday	M
Tuesday	T
Wednesday	W
Thursday	H
Friday	F
Saturday	S
Sunday	U

To combine days, such as Monday and Wednesday, simply enter the two codes together (MW or WM). (The order of the codes is not significant to the scheduler, though you may want to enter the codes in a specific order to help you when reviewing your schedule at some later time.)

Press F2 to save the information entered into this screen.

## Scheduling a Weekly Process

For a given process, when prompted for Reschedule Frequency, enter "W". The following screen will appear.



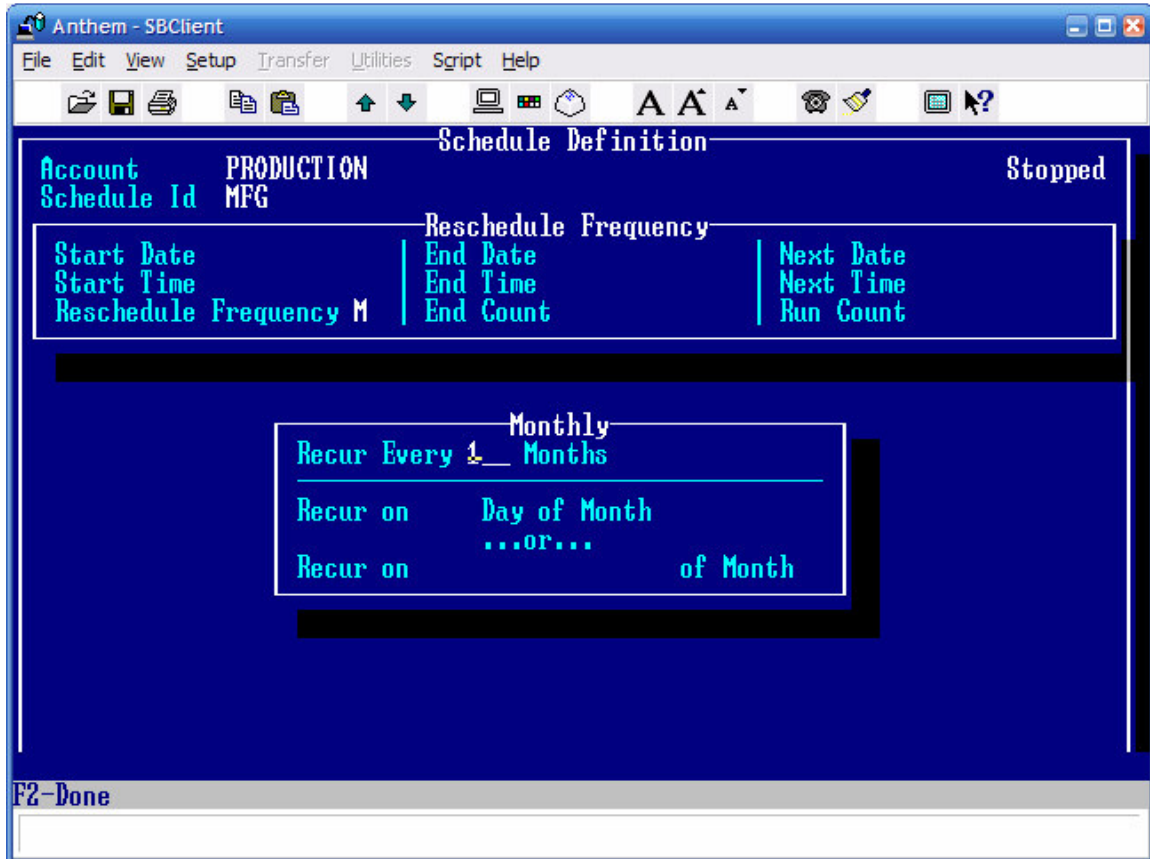
At the first prompt, enter the number of weeks between runs. If you want to run this every week, enter 1. To run the process every other week, enter 2. If you want the job to run only on certain days in that week, enter those days (using the codes previously described) into the Days of Week field.

For example, to run the process every other week, but only on Wednesday and Friday, enter 2 at the first prompt and WF (or FW) at the Days of Week prompt.

Press F2 to save the information entered into this screen.

## Scheduling a Monthly Process

For a given process, when prompted for Reschedule Frequency, enter "M". The following screen will appear.



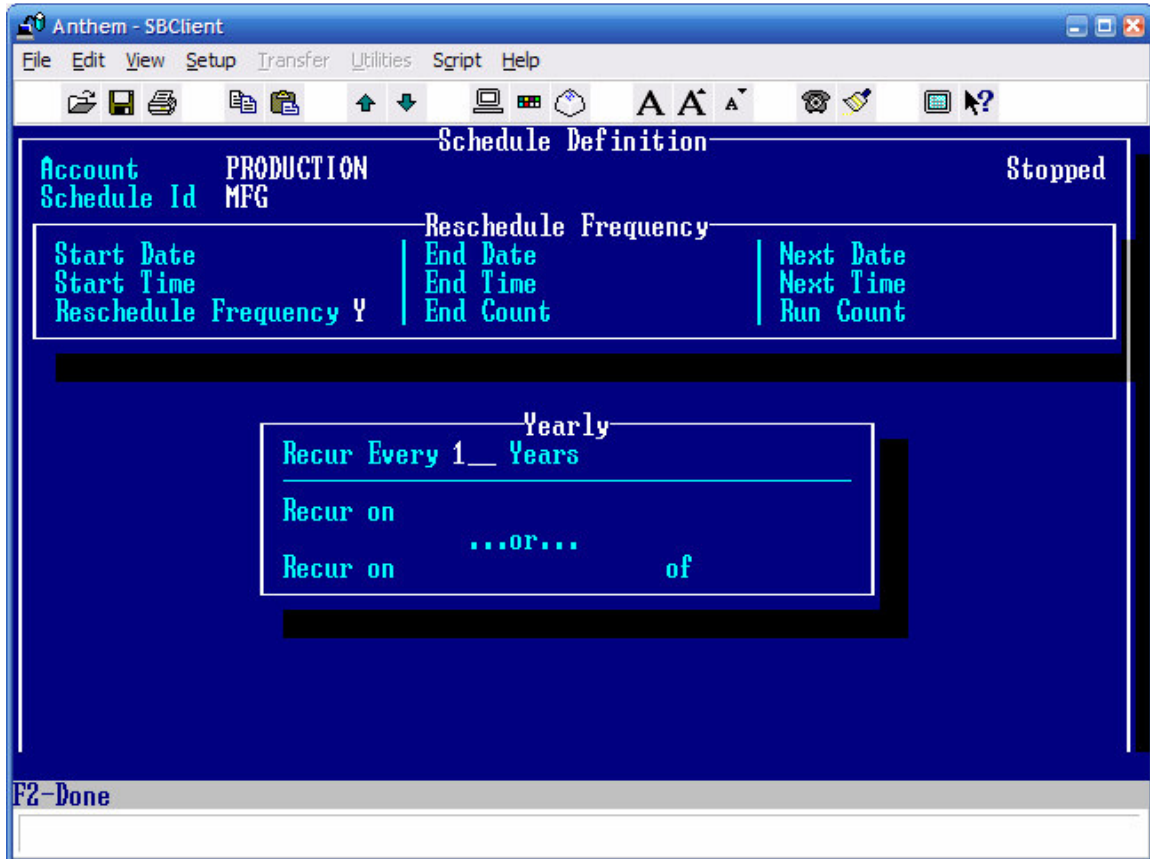
At the first prompt, enter the number of months between runs. To run this process every month, enter 1. To run this process every other month, enter 2. To run the process quarterly, enter 3.

Below the line, you may choose from one of two options for monthly scheduling. The job can either be run on a specific day (by day number) or a specific day based on its position in the month. For example, to run the job on the 15<sup>th</sup>, enter 15 at the first prompt. Or, if you want to run the job on the second Tuesday, leave this first prompt blank and at the Recur On prompt, enter 2 (second), then T (Tuesday). F3 is available on all these prompts to guide you. (You may choose from the first, second, third, fourth, or last occurrence of a given day in a month, and the scheduler will calculate the appropriate day from this information.)

Press F2 to save the information entered into this screen.

## Scheduling a Yearly Process

For a given process, when prompted for Reschedule Frequency, enter "Y". The following screen will appear.



At the first prompt, enter the number of years between runs. To run this process every year, enter 1. To run this process every other year, enter 2.

Below the line, you may choose from one of two options for yearly scheduling. The job can either be run on a specific day in the year or a specific day based on it's position in the year. For example, to run the job on September 15<sup>th</sup>, enter 9/15 at the first prompt. (Ignore the year that appears, that's for reference only and is stripped off when the scheduler calculates the next year.) Alternatively, if you want to run the job on the second Tuesday of September, leave this first prompt blank and at the Recur On prompt, enter 2 (second), then T (Tuesday), then 9 (September). F3 is available on all these prompts to guide you. (You may choose from the first, second, third, fourth, or last occurrence of a given day in a month, and the scheduler will calculate the appropriate day from this information.)

## Scheduling a Once Only Process

A "once only" process will run once at the appointed time (or immediately if no start date/time is entered) and then not be rescheduled. Enter "O" for the Reschedule Frequency to elect a once only process. No additional information is required for this kind of a scheduled item.

## **"Odd" Schedules**

There may arise a situation where you need to schedule a process to run on a schedule that is not directly supported by the Schedule Definition options. When this occurs, configure your process to run on the schedule that includes all of the times you want and possibly more. Then use a wrapper process to interface to the original process and filter out the few instances you might not want.

For example, consider a process to be executed every Wednesday, but not if the Wednesday is the 15<sup>th</sup> of the month. To accomplish this, you might create a wrapper process that verifies if the day is the 15<sup>th</sup> of the month and if it is, the process does not get called. Something like this SB+ paragraph perhaps:

```
IF @DATE"DD" # 15 THEN
  EXEC @PARAM
END
```

Then instead of calling your process directly through the scheduler, call your process via this wrapper on a weekly schedule. For example, if your process is called DO.STUFF and this wrapper is called FILTER.15, the process name in the scheduler would be:

```
FILTER.15,DO.STUFF
```

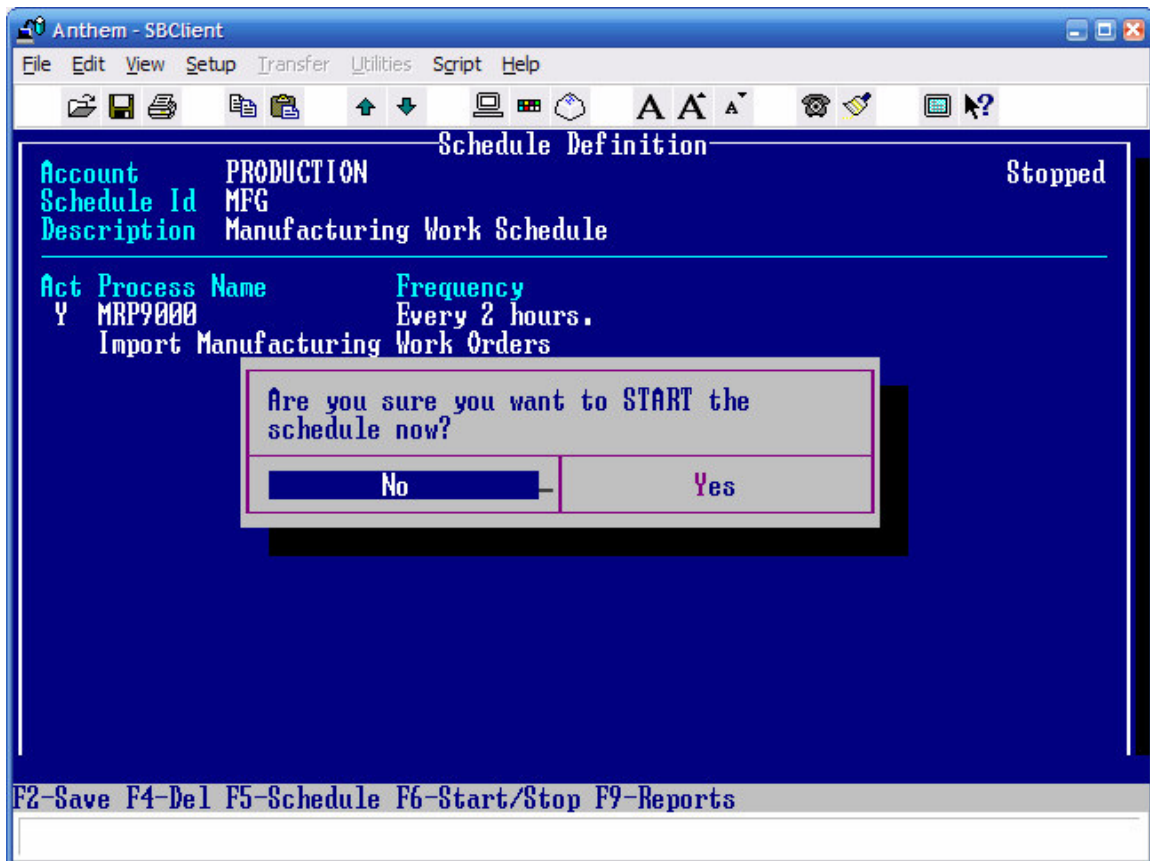
With this in place, the scheduler will invoke the process FILTER.15 on every Wednesday. FILTER.15 then will verify the day number and if not the 15<sup>th</sup>, will invoke the process you originally intended.

## Starting/Stopping the Scheduler

From the main schedule definition screen, you can start or stop the execution of the scheduler. When the scheduler is started, a background process is launched to invoke the scheduled processes. When the scheduler is stopped, this background process is terminated.

In the upper right corner of the main screen is an indicator that describes if the scheduler is running. When the scheduler is running, this indicator will tell you the effective port number. If the scheduler is stopped, the word "Stopped" will be displayed here.

To start a schedule, press the F6-Start/Stop key from the main screen. The following message will be displayed:



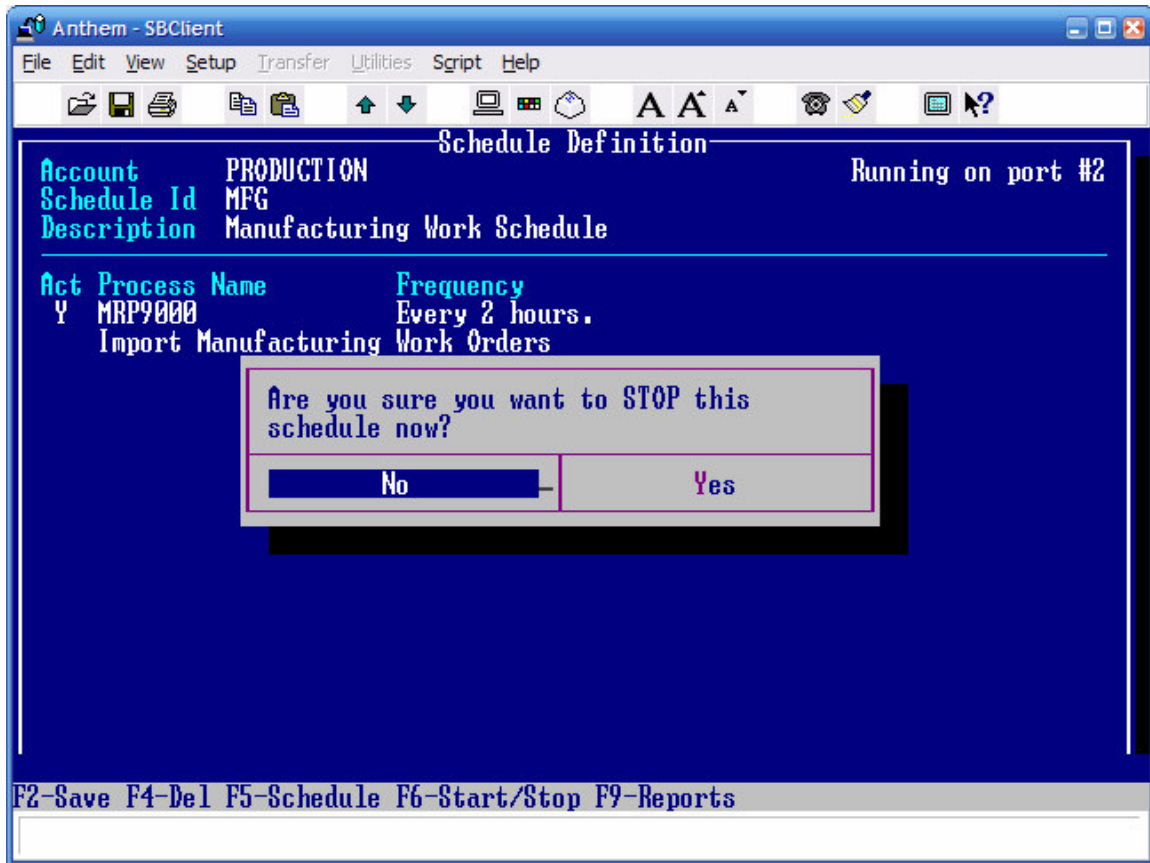
If you want the schedule to be started, press Y. If this key was pressed accidentally, press N or <cr> to proceed without starting the schedule.

Once started, the scheduler will read the active items from the schedule and will proceed to run the processes at their appointed times. Note that if a schedule has no active items, the scheduler will run but nothing will be executed until an active item is added to the schedule or an inactive item is re-enabled.

It's important to press F2-Save or <Escape> after the scheduler has been started so that you release the record lock being held by the entry screen. The scheduler is unable to start any jobs while the screen is locking the schedule details (that is, while the schedule is being edited).

Also note that once the scheduler has been started, you may log off your current session with affecting your scheduled processes. The scheduler will ensure that your scheduled processes will continue running on a background session.

If the scheduler is already running for this schedule and the F6-Start/Stop button is pressed, you will be allowed to stop it. The following message is displayed:



Press Y to stop the scheduler. Or, if this key was pressed by accident, press N or <cr> to proceed without stopping the scheduler. Note that when the scheduler is stopped, the processes listed on that particular schedule will not be executed at their appointed times.

## Viewing Reports

From the main screen, press F9 to see a list of reports available. Move the cursor with the arrow keys and select the report to view with the <cr> key. Each report asks for information to use in selecting the records to view; enter the appropriate values to filter the information shown on the listing.

The following standard reports are available:

**Scheduler Report** – A generic list of all processes that have been launched by the scheduler.

**Scheduled Processes Currently Running** – A list of processes that have been launched by the scheduler and have not yet completed.

**Scheduled Processes That Have Finished** – A list of processes that have been launched by the scheduler and have run to completion.

**Scheduler Status Report** – A list of schedules defined by the /PJSD tool and their status (running, stopped, etc.)

## The Scheduler

The core of the Precision Schedule Definition tool is the scheduler itself. This routine runs invisibly as a background process and accepts input from the /PJSD tool. Note that there will be one background process per schedule as defined by this tool.

The scheduler is effectively an "envelope" around the standard SB+ Job Schedule tool. When the scheduler identifies a process to be run, it wraps a Job Schedule Definition around the process to be run and then invokes that job schedule as a background process. Because the actual scheduler process runs as a normal SB+ Job Schedule process, there are a couple of important details you should be aware of.

First, if the scheduler invokes a job schedule process defined to run in the background, the second job schedule may be invoked as a normal process call on the same port and not on its own independent background process. Therefore, the scheduler will be delayed from starting another process while it is running a process. Furthermore, any other processes that come "due" during the time the other program is running may not be started in a timely manner. Note that this anomaly is based on your system's configuration and may or may not be an issue in your particular environment.

It is also imperative to note that if the process running as a background process prompts for input, the phantom process – **and the scheduler running that phantom process, if on the same port** – will terminate. You can look in the \_PH\_ record (Unidata) or &PH& record (Universe) for more information about the cause of the failure.

Another important thing to keep in mind is that due to record locking restrictions, the scheduler will not run any processes while the schedule is being edited using the /PJSD tool. Therefore, strive to spend as little time as possible editing in the /PJSD tool while the schedule is running.

You can also incorporate the scheduler into your own programming via the SCHED0037 interface process. To use this process, execute the SCHED0037 process passing the name of your schedule as a parameter. This will start the schedule running in the background just as if it had been started from the /PJSD tool. For example:

```
EXEC 'SCHED0037,MYSCHEDULE'
```

...will start the scheduler for the schedule named "MYSCHEDULE".

## The Run Log

When the scheduler starts a process, it first writes a record to the run log, a file named PSI.RUN.HST. In this file you'll find a variety of details about the running job including:

- Schedule Name
- Process Name
- Actual Start Date
- Actual Start Time
- Actual End Date
- Actual End Time
- Next Scheduled Date
- Next Scheduled Time
- User Id (who started the scheduler)
- Port Number
- Account Name
- Status (L=Launching, R=Running, F=Finished)
- Messages (with date and time) from the scheduler.

You can also add your own messages to the run log for a given program instance. In your program, call the process SCHED.LOG passing the message as a parameter. For example, to log the message "I'm done now", the following line might be added to a paragraph:

```
EXEC "SCHED.LOG,,I'm done now"
```

This message will then appear on reports of completed jobs.

Note the two commas. Between these two commas you can place the name of a process to be called to update your own custom attributes in the PSI.RUN.HST record (which will be in @OTHER.REC). You can then update those attributes at the same time as logging a message. (If the message parameter is omitted you can update the run history record via your custom process without adding a message.)

## Frequently Asked Questions

**Q: Can /PJSD be used to run existing job schedules at regular intervals?**

A: YES! However, note that SB+ runs all subordinate job schedules on the original port, so a second background port will not be used, regardless of the Background (Y/N) setting in the job schedule definition.

**Q: Can /PJSD be used to run other types of processes rather than just Job Schedule processes?**

A: Yes, however, the process to be run must not prompt for input, or if it does prompt for input there must be sufficient data stacked by the program to fill in all of the prompts.

**Q: If I stop the scheduler while it's running a job, will it terminate that job?**

A: No, the scheduler will stop only between jobs.

**Q: Sometimes when I stop the scheduler, it's very fast, other times it's very slow. Why the difference?**

A: If the background process has terminated unexpectedly, as happens when one of the processes in the schedule prompts for input, the /PJSD tool will send a request to the scheduler to stop, but the scheduler is already stopped! The delay – about a minute – let's the /PJSD tool figure out that the scheduler is probably no longer running.

Note that the schedule may not actually be stopped, but it may be in the middle of a big time-consuming process. The /PJSD screen will then show that the scheduler has stopped, but the scheduler will continue to run until the current job is completed before terminating.

**Q: How can I figure out the date and time a specific job was started?**

A: The scheduler updates a record in the file PSI.RUN.HST when a process is started. In this record you can find a variety of information including the schedule name, process name, start date/time, end date/time, next date/time, user ID, status, and any messages that may have been generated by the scheduler or running process(es).

**Q: How accurate is the scheduler? Will my processes always get started at the right time?**

A: If not delayed by running a process or the schedule being edited in the /PJSD screen, the scheduler will start a job usually within 10 seconds of its scheduled time. If the schedule is being edited by the /PJSD screen, the job will be started after the changes have been saved. If the scheduler is delayed by running another job, the scheduled task will be started within about 10 seconds of the completion of the first job.

**Q: On Friday I scheduled a job for Mondays only, and when I started the scheduler (on Friday) the process was started. Should it not have waited until Monday?**

A: If the process should not run the first time until Monday, then Monday should be listed as the start date in the schedule. If there is no start date on an item, that item will be run immediately when the scheduler is started. After that first run, the process will adhere to the rules of the schedule.

**Q: Can I start a schedule from my program instead of going through the /PJSD tool?**

A: Yes, invoke the process SCHED0037 passing the name of the schedule to be started as the parameter. For example, the following might be used in a paragraph: EXEC 'SCHED0037, *scheduleName*'

**Q: Didn't this product originally go by another name?**

A: Yes, but due to hyper-aggressive trademark attorneys and their clients who are more interested in threatening good people rather than actually solving business problems, the name was changed on July 14, 2006.

## **Precision Solutions *Empowerment* Initiative**

In late 2003, Precision Solutions launched a new approach to software development. Titled the *Empowerment Initiative*, this infrastructure provides a foundation for solutions that customers can easily adapt and enhance over time.

Rather than deliver product with restricted features for modifications and enhancements, Precision Solutions software is built to empower the customer to adapt the software to their own particular environment and needs. To this end, all software is delivered with source code and documentation, the code is written to exacting standards of accuracy, compatibility, understandability, and expandability, and training is available to assist the customer in expanding the product for their own particular needs.

Precision Solutions is pleased to add the Schedule Definition tool to our family of *Empowerment*-enabled development products.