

# **IBM MVS System Trace Analyzer User Guide**

March 2011

IBM T.J. Watson Research

## Contents

### [IBM MVS System Trace Analyzer User Guide](#)

- [Getting Started](#)
  - [Introduction](#)
  - [Installing Graphviz](#)
  - [Tutorial](#)
  - [HASID](#)
- [Shortcuts for Navigating and Exporting Traces](#)
- [Tasks](#)
  - [Saving a Subset of the Events to a New Trace File](#)
  - [Find events in Events Flow View](#)
  - [Make event selection in the Events Flow View and statistics views](#)
  - [Manage component definitions for events](#)
  - [Manage default search terms](#)
  - [View event sequence in sequence context graph](#)
  - [View statistics in Event Type Statistics, Location Statistics, and Pathology Report view](#)
  - [Collecting a Module Map using Hardware Instrumentation Services on z/OS](#)
- [Reference](#)
  - [Events Flow View](#)
  - [Statistics views](#)
    - [Event Type Statistics view](#)
    - [Location Statistics view](#)
    - [Pathology Report view](#)
    - [Sequence Context view](#)

## Introduction

System z provides a number of ways to create traces, such as using the IPCS SYSTRACE command. Traces from this utility can describe branches or other events of interest, with their corresponding HASIDs, work units, addresses, modules or other information. This information can be used for debugging or performance analysis. However, the challenge is often the sheer size of these traces. It is not unusual to get traces with millions of lines. Understanding and navigating these textual traces can be a daunting task, even with good text searching tools.

IBM MVS System Trace Analyzer is a visual tool that allows you to easily explore, navigate and understand these traces. It has the following features:

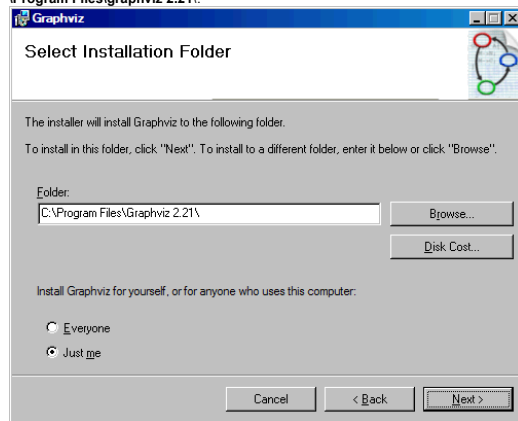
- IBM MVS System Trace Analyzer can visualize large traces (millions of lines) efficiently, at a detailed level, as well as at an overview level.
- Patterns in the trace are easy to spot visually.
- You can easily see if there is activity in nucleus versus user regions.
- You can easily navigate traces with fast zoom and pan functions, as well as powerful search functions.
- IBM MVS System Trace Analyzer shows scheduling information such as HASID, work unit or processor, and allows you to filter based on these parameters.
- You can immediately drill down on "debug events" (abends) and show the relevant preceding events (for example, in the same HASID or work unit)
- IBM MVS System Trace Analyzer can show "diagnostic events" that you created programmatically in the trace.
- IBM MVS System Trace Analyzer helps with performance analysis by showing repetitive patterns, context switches and spinning events.
- IBM MVS System Trace Analyzer provides a seamless transition between a textual and a graphical representation of the trace file. No information is discarded or withheld.
- IBM MVS System Trace Analyzer uses a unique system of "murals" on each side of the canvas from which you can see the context, even if zoomed in deeply.
- IBM MVS System Trace Analyzer can help to reverse-engineer legacy applications by showing the flow of execution in a natural way.
- You can isolate and save smaller areas of interest in the trace, for later inspection or to share with other users.
- Multiple traces can be loaded simultaneously in IBM MVS System Trace Analyzer.

## Installing Graphviz

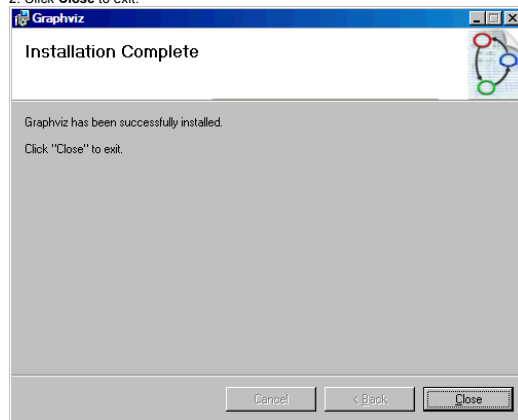
Installing the Graphviz executable is a prerequisite for using the **Sequence Context** view. You can download the Graphviz installation package from [http://www.graphviz.org/Download\\_windows.php](http://www.graphviz.org/Download_windows.php), and use the latest stable version to install Graphviz.

Follow these steps to install the Graphviz executable on your computer:

1. Run the Graphviz installation directory on the following wizard page. Then click **Next**. For this example, the installation directory is specified as **C:\Program Files\graphviz 2.21**.



2. Click **Close** to exit.



3. Update the system path.

- For Windows, you need to update the system path by entering the **Control Panel**, and select **System->Advanced->Environment Variables**.
- Select **Path** and add **C:\Program Files\graphviz 2.21\bin** to the path (here the path is specified as **C:\Program Files\graphviz 2.21\bin** because the installation directory is **C:\Program Files\graphviz 2.21** for this example).
- Open a new command prompt window and run the command **dot -V** to verify if the path has been added successfully. If you add it successfully, the window shows the version as follows.

```

Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Administrator>cd \

C:\>dot -U
dot - Graphviz version 2.20.3 (Wed Oct 8 06:02:12 UTC 2008)

C:\>_

```

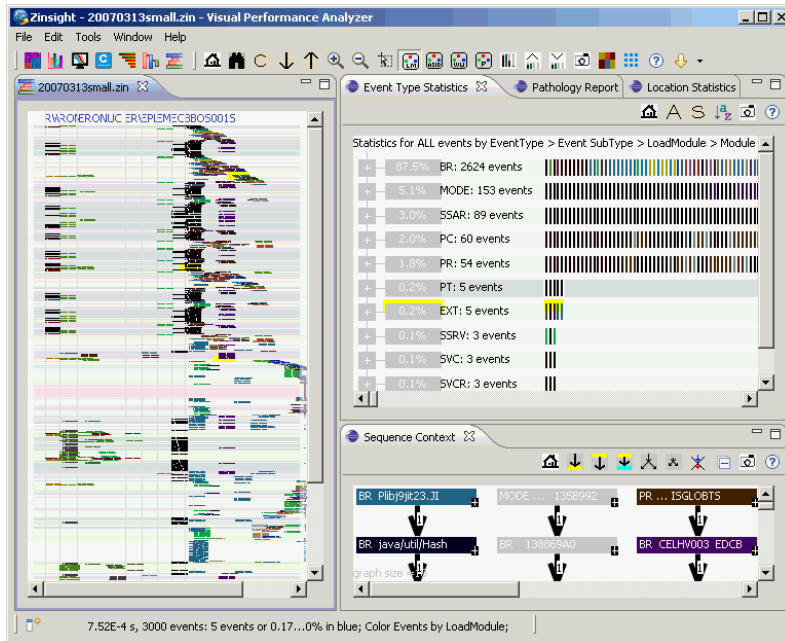
leftmargin="0" topmargin="0" style="background-color: white;" lang="EN-US" link="blue" marginheight="0" marginwidth="0" vlink="blue">
**Getting started**

You can open a trace file with the .zin or .ziz extension from the **Open File** dialog by selecting **File -> Open file**.

IBM MVS System Trace Analyzer helps the user to discover the right information in these traces trace with 3 views:

1. The Events Flow view allows the user to explore the trace in a visual rather than textual way. (shown on the left in the figure below)
2. The Statistics views (the Event Type Statistics view, the Pathology view and the Location view) show the events by type and in descending order of frequency. (shown on the upper right in the figure below)
3. The Sequence Context View can show patterns in the execution to answer questions like: "How did I get to a certain type of events", or "what is the context of a hotspot". (shown on the lower right in the figure below)

In case any of the views are closed or resized and you want to reset the perspective to its defaults, select **Window -> Reset Perspective**.



## HASID

HASID stands for Home Address Space Identifier.

### Address space

The range of virtual addresses that the operating system assigns to a user or separately running program is called an *address space*.

### ASID

In some ways a z/OS address space is like a UNIX process, and the address space identifier (ASID) is like a process ID (PID).

## Keyboard Shortcuts to work with IBM MVS System Trace Analyzer

Zooming and panning:

**Home key** : Show complete trace

**PageUp/Dn** : page up and down

**Arrow keys** : pan left, right, up, down

**+ and -** : zoom in and out

**< and >** : compress and expand vertically

**[ and ]** : compress and expand horizontally

Drag a rectangle with mouse to zoom in on an area

Press the key **Ctrl** while holding the left mouse button to drag the canvas of the Events Flow View or view so as to pan in all directions

**IBM MVS System Trace Analyzer Editor:**

**a** : export all the events in the Events Flow View into a new file  
**c** : clear all highlights in all views  
**d** : show or hide debugging events  
**e** : export highlighted events in the IBM MVS System Trace Analyzer Events Flow view to a new file  
**f** : find events to highlight  
**g** : gray out non-highlighted events (toggle)  
 left-click in scheduling bar to highlight HASID, WU or Proc  
 right-click on any event for more options  
**m** : highlight thread migrations  
**s** : show or hide scheduling  
**t** : show or hide text mode when zoomed in  
**v** : export highlighted events to new file in the current screen of Events Flow View

**A** : color by HASID  
**W** : color by Work Unit  
**P** : color by Processor  
**S** : color by LoadModule and show scheduling  
**L** : color by LoadModule (default)

**Sequence Context view:**

**c** : clear all highlights in all views

**Event Type Statistics, Pathology Report and Location Statistics views:**

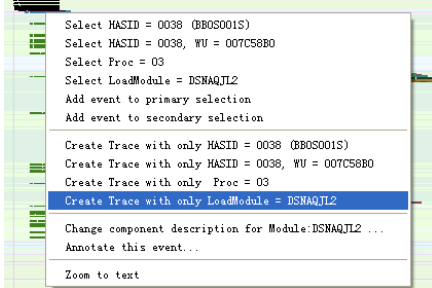
**c** : clear all highlights in all views  
**a** : calculate the statistics for all events  
**s** : calculate the statistics for the highlighted events

### Saving a Subset of the Events to a New Trace File

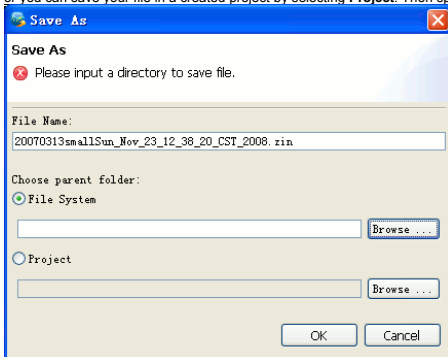
Right-click an event in the Events Flow View, and you can create trace files with specific events in the same HASID, WU, processor or loadmodule.

The following screen capture shows the context menu options for creating a trace file. The second group shows the options for creating trace files. You can create trace of different event groups as follows:

- Select **Create Trace with only HASID = HASID** to create trace with the events in the same HASID.
- Select **Create Trace with only HASID = HASID, WU = WU** to create trace with the events in the same work unit.
- Select **Create Trace with only Proc = processor** to create trace with the events on the same processor.
- Select **Create Trace with only LoadModule = loadmodule** to create trace with the events in the same loadmodule.



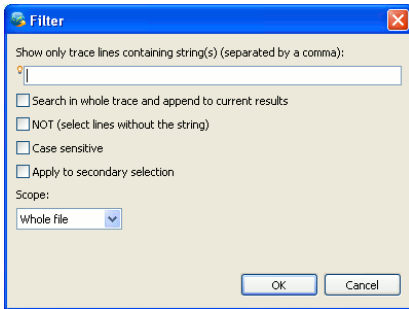
After you selected the option in the preceding screen capture, a dialog opens for you as follows. You can choose to save the file on your local or remote file system by selecting **File System**, or you can save your file in a created project by selecting **Project**. Then specify a directory for the file system or project, and click **OK**.



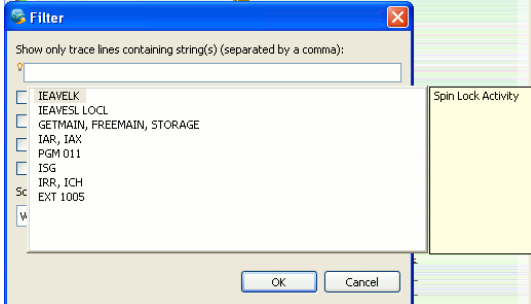
Then the events of the same attribute are exported to a new file which is loaded in the Events Flow View.

### Find events in the Events Flow View

To find events in the Events Flow View, press the shortcut key **F**, or click the button  on the IBM MVS System Trace Analyzer toolbar to open the **Filter** dialog as follows.



In the field of this dialog, type any attribution(HASID, WU, loadmodule...) of the events to find specific events in Events Flow View. You can type several attributions and each is separated by a comma. You can also press **Alt + /** to get the default search terms and select one attribution from the term list to fill in the field, as shown in the following screen capture. For each attribution in the search term list, there is a tag containing the description for this attribution. You can manage the default search terms in the **Preferences** dialog(see [Manage default search terms](#)).




Some search conditions are also available for you:

- **Search in whole trace and append to current results**  
Append the current search results to the search results of last time, which means the search results of last time are still highlighted.
- **NOT**  
Highlight the event lines that don't contain the search string.
- **Case sensitive**  
Search the event lines with case-sensitive search strings.
- **Apply to secondary selection**  
Add your current search results to the secondary selection, which is highlighted in turquoise.
- **Search scope**  
The search scope can be the whole file or the current results. If you select **current results**, you can search for a second time within the search results currently highlighted.

The search results will be highlighted in yellow color in the Events Flow View and Statistics Views.

## Make event selection in the Events Flow View and statistics views


Views provides you with many ways to select the events you want. You can perform a single event selection, multiple event selection, cross-view selection, and so on.

With the button  enabled, you can select events in the Events Flow View.

To select a single event, you can simply click the event (shown as a rectangle).


To select multiple events, there are several ways for you:

- Select the event you want with the **Shift** key pressed. For each time click a single event.
- Press the **Shift** key, and drag an area of events in Events Flow View with your mouse, and then a group of events are selected and highlighted in yellow.
- After you select an event, right-click another event that is not highlighted and select **Add event to primary selection**. You can add multiple events to primary selection with the context menu option.
- Click on the vertical bars on the left side to select a group of events. The following screen capture shows the result effect.




By clicking the tree nodes or graph nodes, you can select event nodes in the statistics views. The **Shift** key also applies to the statistics views and the selection across views.

To select a group of events of the same attribute, such as HASID, WU(work unit), processor, and loadmodule, right-click an event and select corresponding options, as shown in the following screen capture.



To select two groups of events, right-click any event in Events Flow View, and select **Add event to the primary selection** as one group, or select **Add event to the secondary selection** as the secondary group. The primary group of events are highlighted in yellow and the secondary group of events are highlighted in turquoise, as shown in the following screen capture.



After you select a group of events, you can remove one or more events from the group by right-clicking a highlighted event and select **Remove event from primary selection** or **Remove event from secondary selection**, depending on the group this event belongs to.

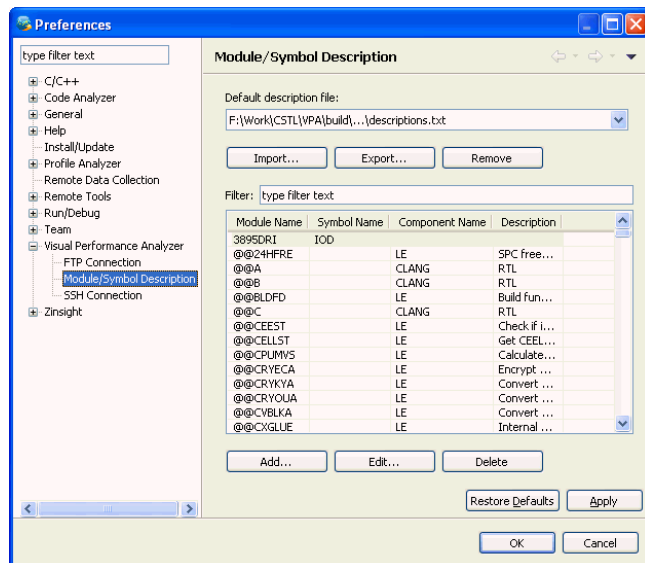
**Note:** By clicking the event in any of the views, the other views highlight the corresponding event.

## Manage component definitions for events

You can manage the description for the modules, symbols, and components in the **Module/Symbol Description** panel.

To open the **Module/Symbol Description** panel, perform these steps:

1. Click **Window -> Preferences**. The **Preferences** window is displayed.
2. In the **Preferences** window, click **Visual Performance Analyzer -> Module/Symbol Description**.



### Import a description file

1. Click the **Import** button.
2. In the **Import a description file** window, select the description file to import, and click **Open**. After you import a valid description file, all the names of the modules, symbols, and components, and their descriptions are displayed in the **Module/Symbol Description** panel. The imported file is in the **Default description file** list.
3. Click **Apply** or **OK**.

### Export a description file

1. Select a description file in the **Default description file** drop-down list and click **Export**.
2. In the **Export a description file** window, specify the directory and file name to save the exported file, and click **Save**.

### Remove a description file

You can remove a description file from the **Description file** drop-down list:

1. Select a description file in the **Description file** drop-down list.
2. Click **Remove**. In the message window that is displayed, specify whether to remove the current description file. Click **Yes** in the message window.
3. Click **Apply** or **OK**.

### Change a default description file

To choose a description file to be the default description file, follow these steps:

1. Select a description file from the **Default description file** drop-down list.
2. Click **Apply** or **OK**.

### Add, edit, or delete a line in the description file

**To add a line to the file, do the following steps:**

1. Select a file you want to edit in the **Description file** drop-down list.
2. Click the **Add** button. The **Add a description** window is displayed.
3. Specify or select a module name in the **Module Name** combo box. This field is required.
4. Optional: Specify a symbol name in the **Symbol Name** field.
5. Optional: Specify a component name in the **Component Name** field.
6. Specify the description for the module or symbol in the **Description** text area. The description is for the module if you leave the **Symbol Name** field empty; otherwise the description is for the symbol if you specify a name in the **Symbol Name** field.  
**Note:** Do not type any tabs or carriage returns in the description.
7. Click **OK** in the **Add a description** window.
8. In the **Preferences** window, click **Apply** or **OK**.

**To edit a line in the file, do the following steps:**

1. Select a file you want to edit in the **Description file** drop-down list.
2. Select a line in the preview table. The following figure shows an example.
3. Double-click the selected line or click the **Edit** button. The **Edit Module/Symbol Description** window is displayed.

- Specify the names and description in the fields.
- Click **OK** in the **Edit Module/Symbol Description** window.
- In the **Preferences** window, click **Apply** or **OK**.

To delete a line from the file, do the following steps:

- Select a file in the **Description file** drop-down list.
- Select a line in the preview table.
- Click **Delete**.
- Click **Apply** or **OK**.

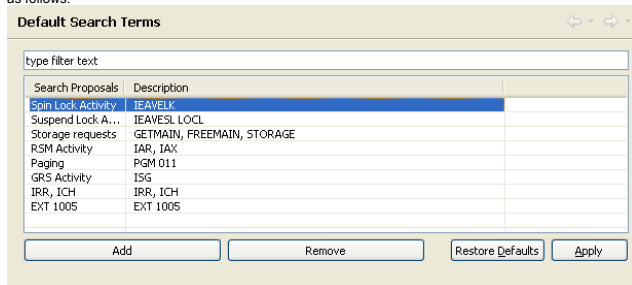
#### Reset descriptions to default

To restore default settings, click the **Restore Defaults** button, and then Click the button **OK** or **Apply** to confirm the settings. Then all the settings you set is removed.

### Manage default search terms

When you search the events with the filter dialog, you might use the default search terms to help you search an event. The default search terms include a search string and the search proposal for this string.

To manage the default search terms, open the editing page of the default search terms by selecting **Window->Preferences->Default Search Terms**. The **Default Search Terms** page is open as follows.



You can define the search strings in the **Description** column and the search proposal for each search string in the **Search Proposals** column.

To edit a search term, simply edit the terms in the table, and then click **Apply** and **OK** to save the changes.

To add a search term, click **Add** and edit the new line. Then click **Apply** and **OK**.

To remove a search term, select one line and click **Remove**. Then click **Apply** and **OK**.

### View event sequence in sequence context graph

In the **Sequence Context** view, you can view the events selected in the Events Flow View or any other statistics views in the form of a forward or backward sequence context graph.

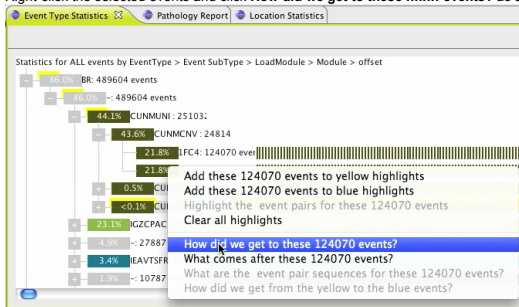
In the progress of generating the sequence context graph, you can terminate it by clicking **Cancel**.

**Prerequisite:** Before you work with the **Sequence Context** view, you must install Graphviz on your computer(see [Installation Graphviz](#)).

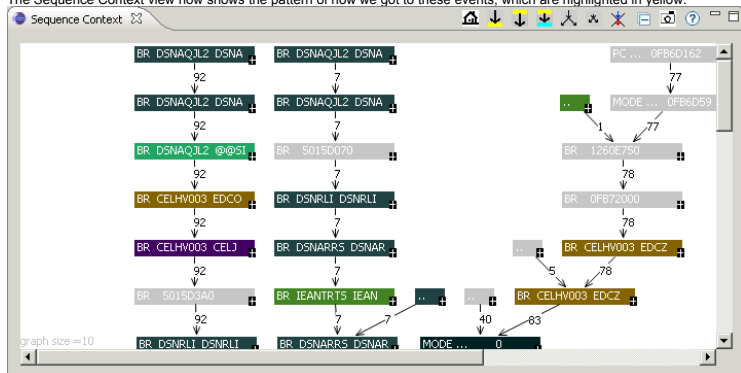
#### How did we get to a certain type of events

To build a backward sequence context graph, select one or multiple events in the Events Flow View or the statistics views, and do one of the following actions:

- Click the button on the toolbar of the **Sequence Context** view.
- Press the shortcut key **b** after you activate the **Sequence Context** view.
- Right-click the selected events and click **How did we get to these nnnn events?** as shown in the following figure.



The Sequence Context view now shows the pattern of how we got to these events, which are highlighted in yellow.



#### What comes after the events

To build a forward sequence context graph, select one or multiple events in the Events Flow View, and do one of the following actions:

- Click the button on the toolbar of the **Sequence Context** view.



- Press the shortcut key **f** after you activate the **Sequence Context** view.
- Right-click the selected events and click **What comes after these nnnn events?**


### How to get the event pair sequences

Sequence Context view can help with the analysis of paired events.

In the Events Flow View or the statistics views, select one or multiple events, right-click and select **What are the xxxx event pair sequences for these nnnn events?**

### How did we get from a type of events to another type of events


To build a sequence context graph from a certain type of events to another type of events:

1. Select one or multiple events in the Events Flow View, right-click and select **Add event to yellow highlights**.
2. Select one or multiple events in the Events Flow View, right-click and select **Add event to blue highlights**.
3. and do one of the following actions:
  - Click the button  on the toolbar of the **Sequence Context** view.
  - Press the shortcut key **t** after you activate the **Sequence Context** view.
  - Right-click the selected events and click **How did we get from the yellow to the blue events?**

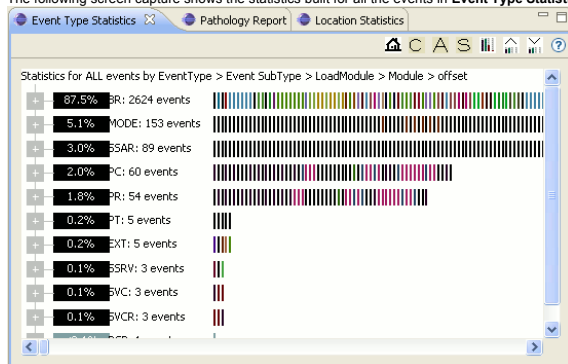
## View statistics in Event Type Statistics, Location Statistics, and Pathology Report view

You can build statistics for all or selected events in the tree views: **Event Type Statistics**, **Location Statistics**, and **Pathology Report** view.


### Build statistics for all the events

To build statistics for all the events in any of the three views, click the button  on the toolbar of the view accordingly. By default, the statistics are built for all the events in these views after you load your .zin or .ziz file.

The following screen capture shows the statistics built for all the events in **Event Type Statistics** view.



### Build statistics for the selected events


To build statistics for the selected events in any of the three views, select any events you want in the views, and then click the button  on the toolbar of the view accordingly.

You have several ways to make your selection:

1. In the Events Flow View, find and select one or more events (press the **Shift** key for multiple selection).
2. In the **Sequence Context**, select one or more nodes.
3. In the **Event Type Statistics**, **Location Statistics**, or **Pathology Report** view, select one or more tree nodes. You can also select the event by clicking on the event bars.

### View matching events

In a trace file, some events appear in matched pairs. For examples, SVC events are matched with their corresponding SVCR events. You can find matching events pairs by right-clicking on a node in the statistics view and choose **Select Matching Events**, as shown in the following picture.




Then the matching events are highlighted. The previously selected events are highlighted in yellow and the matched events are highlighted in turquoise correspondingly. The following picture shows this effect.




By hovering over the event bars, you can get the information of the elapsed time in the ToolTip.

### View elapsed time for matching events

By default, the elapsed time of matching events is hidden when you open the statistics views. You can view them by clicking the button  on the toolbar of statistics views, and then you will find that all the matching events are highlighted in light gray, which is shown in the following screen capture.



By clicking the toolbar button , you can increase the bar height for matching events, and you can compare the proportion of the time from the event-bar height by zooming in to a certain extent.



## Collecting a Module Map using Hardware Instrumentation Services on z/OS

### Overview

On z/OS V1R8 or later, Hardware Instrumentation Services (HIS) is available to capture a module map showing the virtual address ranges of the various modules loaded on your system. This

document describes how to capture this module map without using other features available in HIS.

This document describes how to collect a z/OS Module Map in 3 steps:

**Step 1** [Verify System and Software Requirements](#)

**Step 2** [Configure z/OS to run HIS and capture the module map](#)

**Step 3** [Collect the Module Map on z/OS](#)

### **Step 1** Verify System and Software Requirements

All the following requirements must be met in order to successfully capture CPU MF Sample data.

#### Software:

The z/OS LPAR being measured must be at z/OS 1.8 or higher and PTFs for the following APARs must be applied.

- z/OS 1.8 - APARs OA25755, OA25750, and OA25773
- z/OS 1.9 - APARs OA25755, OA25750, OA25773, and OA32113
- z/OS 1.10 - APARs OA25755, OA25773, OA32113, and OA34485
- z/OS 1.11 - APARs OA30429 and OA34485
- z/OS 1.12 - APAR OA34485

### **Step 2** Configure z/OS to collect CPU MF Sample Data

1. To use Hardware Instrumentation Services, you must be running z/OS 1.8 or higher with the PTFs for all the APARs listed in [Step 1, Verify System and Software Requirements](#) above.
  - Note that installing these updates may require an IPL with CLPA.
  - IBM recommends that you check regularly for additional service updates for HIS.
- 2.
3. If you don't have a HIS member in SYS1.PROCLIB, copy the following procedure into the member, HIS, in your PROCLIB. This provides a method to start Hardware Instrumentation Services (HIS).

#### **SYS1.PROCLIB(HIS)**

```
4. //HIS PROC
   //HIS EXEC PGM=HISINIT,REGION=0K,TIME=NOLIMIT
   //CMDFILE1 DD DUMMY
   //CMDFILE2 DD DUMMY
   //SYSPRINT DD SYSOUT=*
```

- 5.
6. Authorize the HIS procedure to write data to the OMVS File System
  - Add a RACF userid for the HIS procedure

#### **TSO Command to add a HIS userid**

```
• adduser HIS omvs(uid(123) home(/HIS*))
```

where UID(123) is the OMVS uid and /HIS is the default home directory.

- **Note:** While OMVS access is required, there is no special authorization needed. Also, any directory can be used for the HOME directory.

#### **Example: Define a HIS userid for the HIS-started task**

```
• //DAEMONS EXEC PGM=IKJEFT01
  //SYSPRINT DD SYSOUT=*
  //SYSTSIN DD *
  ADDUSER HIS OMVS(123) HOME(/HIS*)
  /*
```

- 7.
8. Define the HIS-started task to RACF:
9. To set up the HIS-started task to RACF, you need to define a profile for it to the RACF generic resource class called STARTED using the RDEFINE command.
- 10.

11. **Note:** If the STARTED class is not active, RACLSTed, and GENERIC profile checking is not activated, use the RACF SETROPTS command to activate (CLASSACT), RACLIST, and GENERIC the STARTED class first before you can define a new profile to it. In most environments, this might already have been done. Therefore, you might not have to include the SETROPTS command to CLASSACT, RACLIST, and GENERIC the STARTED class. In this example, we include these SETROPTS commands for completeness. Running these commands when the STARTED class is already activated has no effect.

- 12.
13. **Example: RACF commands to define HIS-started task to RACF**

```
14. //DAEMONS EXEC PGM=IKJEFT01
   //SYSPRINT DD SYSOUT=*
   //SYSTSIN DD *
   SETROPTS CLASSACT(STARTED)
   SETROPTS RACLIST(STARTED)
   SETROPTS GENERIC(STARTED)
   RDEFINE STARTED HIS.* STDATA(USER(HIS) TRUSTED(YES))
   SETROPTS RACLIST(STARTED) REFRESH
   SETROPTS GENERIC(STARTED) REFRESH
   /*
```

- 15.
16. Specify where to store the HIS output file:  
Create the HOME directory in a local USS file system by issuing the `mkdir` command under USS. Also assign read/write/exec authority.

```
17. tso omvs
    mkdir /HIS
```

18.

In this example, /HIS will be the default directory where the HIS output file will be stored.



### Collect the Module Map on z/OS

**Note:** The z/OS LPAR being measured must be at z/OS 1.8 or higher with the PTFs for all the APARs listed in [Step 1, Verify System and Software Requirements](#) above.

1. Issue the following z/OS Command to start the Hardware Instrumentation Services (HIS) Address Space :

**System Console**

```
2. s his
```

3.

The following message indicates a successful start:  
**HIS002I HIS INITIALIZATION COMPLETE**

4. At the appropriate time (**remember, this is a static map, a snapshot of the module map when the map is collected**), issue the following z/OS Command to capture the map.

**System Console**

```
5. t his,b,tt='runid',path='/HIS',maponly,mas=all
```

6.

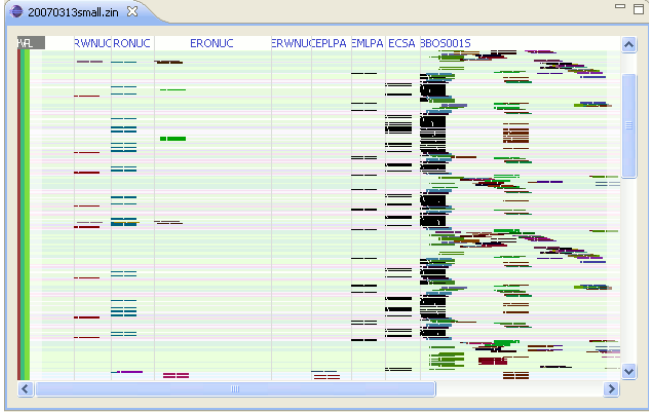
Where:

- **B** tells HIS to Begin collecting the MF Data described hereafter.
- **TT** allows you to specify a unique label or 'runid' to identify this measurement.
- **PATH** identifies where the HIS output files will get written. If the user does not specify the **PATH** option, files will be written to the HIS user's HOME directory specified in the OMVS RACF segment.
- **MAPONLY** tells HIS to immediately capture a module map. No HIS Counter or Sample data will be collected if this option is specified.
- **MAS** or (**MAPASID**) indicates to HIS what address spaces should be included in the map. The **ALL** option (recommended) requests a map of all swapped in address spaces. Optionally, the **MAPASID** or **MAPJOB** option can be used to capture a map of only the address spaces of interest. If you specify **MAPASID (MAS)** or **MAPJOB (MJOB)** you will always map the Nucleus and LPA areas as well.

### Events Flow View

Events Flow View displays the events which constitute the trace file in a graphical form. When you load a .zin or .ziz file, all the events in the trace file are represented as colored rectangular blocks. The placements of the rectangular blocks on x-axis correspond to the address space they belong to, and on the y-axis, the placements are made according to the order of their appearance in the trace file.

The following screen capture displays the Events Flow View.



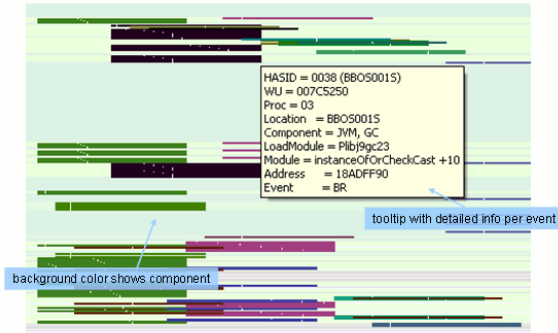
#### The events in Events Flow View

The events are grouped in the regions of different address spaces, which are separated by vertical lines in grey, and the events that belong to the same address space are in the same region. The name of each address space is displayed in blue text on the x-axis on top, as shown in the preceding screen capture.


Each event occupies one line in Events Flow View, and the events are also arranged orderly from top to bottom, which indicates the order of the occurrence of the events.

The events in the same color belongs to the same loadmodule.

Each event has the common attributes: HASID, WU(work unit), processor, component name, location, loadmodule, module, address and event name. When you point the cursor on any of the events, the attributes of the event are displayed in the ToolTip. Each background color in Events Flow View denotes a certain kind of component. The following screen capture gives the illustration of the event in details.



The following screen capture shows a sequence of the events (colored green) zoomed in. There are totally 6 events in the sequence. The white marker in the rectangular block denotes the module of the event, and the location of the marker depends on the address of the module. The text besides the event shows the information of the loadmodule and module of the event.



The following screen capture displays the events in text mode. After you selected **Toggle text mode** in the drop-down list of Events Flow View (see [The menu items of the Events Flow View for Toggle text mode](#)), in which the events are displayed with their attributes shown in lines. Each attribute line corresponds to the line in which a specific event is.


```

03 0038 BBO50015 007C5250 BR 00FF89E8 ROMJLC IGVCPPOOL IGVCPFRE 0
03 0038 BBO50015 007C5250 BR 18133CF8 BBO50015 Pibj9vm23 SnewByteArray 10
03 0038 BBO50015 007C5250 BR 1813E7F0 BBO50015 Pibj9vm23 SnewBaseTypeArray 10
03 0038 BBO50015 007C5250 BR 18100DB8 BBO50015 Pibj9vm23 SZ_NEWBASETYPEARRAY 10
03 0038 BBO50015 007C5250 BR 18A2B780 BBO50015 Pibj9gc23 S39AllocateIndexableObject 10
03 0038 BBO50015 007C5250 BR 18AD7958 BBO50015 Pibj9gc23 SMM_AllocateDescription::allocate(MM_EnvironmentMo
03 0038 BBO50015 007C5250 BR 18ADD778 BBO50015 Pibj9gc23 SMM_EnvironmentBase::unwindExclusiveVMAccess() 1C
03 0038 BBO50015 007C5250 BR 18137CE0 BBO50015 Pibj9vm23 Sj9jni_createLocalRef 10
03 0038 BBO50015 007C5250 BR 181361F0 BBO50015 Pibj9vm23 SexceptionOccurred 10
03 0038 BBO50015 007C5250 BR 1813F0D0 BBO50015 Pibj9vm23 SsetArrayRegion 10
03 0038 BBO50015 007C5250 BR 1810E1D0 BBO50015 Pibj9vm23 SZ_SETARRAYREGION 10
03 0038 BBO50015 007C5250 BR 181361F0 BBO50015 Pibj9vm23 SexceptionOccurred 10
03 0038 BBO50015 007C5250 BR 155A90C8 BBO50015 J Jcom/lbm/ejs/Required.preInvoke(Lcom/lbm/websphere/cs

```

#### Vertical bars

There are three vertical bars placed on the left side of Events Flow View, as shown in the following screen capture.

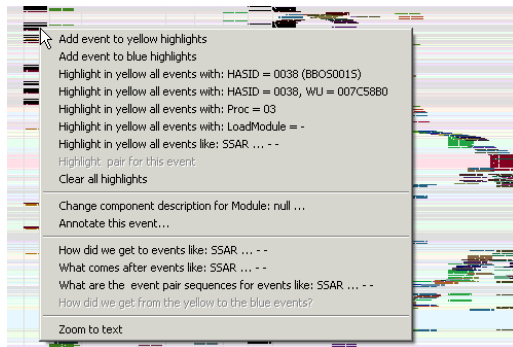


The three bars provides you with quick navigation to spot the events of specific HASID, WU, or Proc(processor) quickly.

- The first bar (the leftmost of the three bars) denotes the HASIDs of the events, and each kind of color on the bar represents a specific HASID that the events belong to. Clicking any part on this bar highlights the events of the same HASID that the color of this clicked part denotes.
- The second bar (the middle one) denotes the work units of the events, and each kind of color on the bar represents a specific work unit that the events belong to. Clicking any part on this bar highlights the events of the same work unit that the color of this clicked part denotes.
- The third bar (the rightmost of the three bars) denotes the processors of the events, and each kind of color on the bar represents a specific processor that the events run with. Clicking any part on this bar highlights the events of the same processor that the color of this clicked part denotes.

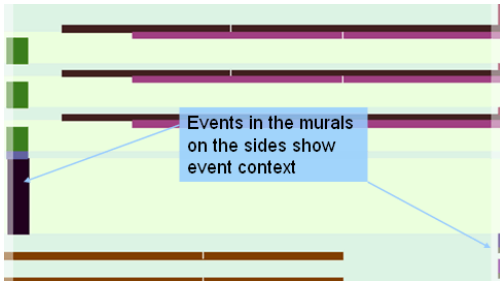
#### The context menu of Events Flow View

Right-clicking any event opens a context menu with some selection options and some options to create trace files derived from the currently loaded file. The context menu of a common event is displayed as follows.



#### The murals of Events Flow View

The murals on the sides show event context when you zoom in, as shown in the following screen capture. The murals give you a perspective of the events which are not displayed in the current window and those events are in the position that is closest to the visible area in Events Flow View. By scrolling horizontally, you can see the "hidden" events.




#### The toolbar buttons of Events Flow View

Button	Name	Description
	Home	Zoom to resize the trace canvas to fit the window.
	Select by pattern	Find and select the events which satisfy certain conditions.
	Clear selected events	Clear all selections made in Events Flow View and all IBM MVS System Trace Analyzer views.
	Pan to the next event	Highlight and draw a line through the next event (by first selecting a group in the statistics views, and click this button to find the events one by one in a forward order).
	Pan to the previous event	Highlight and draw a line through the previous event (by first selecting a group in the statistics views, and click this button to find the events one by one in a backward order).
	Zoom in	Increase the size of the graph.
	Zoom out	Decrease the size of the graph.
	Select Zoom	Select an event, clear the selection made on an event, or drag a rectangular to zoom in.
	Color by LoadModule	Set the events color according to the LoadModule.
	Color by HASID	Set the events color according to the HASID.
	Color by WorkUnit	Set the events color according to the WorkUnit and HASID.
	Color by processor	Set the events color according to the processor.
	Show/hide elapsed time for paired events	Show or hide the elapsed time for paired events.
	Increase bar height for elapsed time	Increase the bar height that denotes the elapsed time of the matching events.
	Decrease bar height for elapsed time	Decrease the bar height that denotes the elapsed time of the matching events.
	Capture screen shot	Capture a screen shot for the current visible trace graph in Events Flow View.
	Hide/show scheduling bar	Hide or display the vertical scheduling bar to the left of Events Flow View.
	Hide/show multiple processors	Hide or display multiple processors in Events Flow View.

#### The menu items of the Events Flow View

The following screen capture shows the drop-down menu of Events Flow View



The drop-down menu list of the Events Flow View contains the following menu items:

- **Save All Events (a):** Export all the events to a different file and load the new file in Events Flow View. 'a' is the keyboard shortcut for the action.
- **Save Highlighted Events Only (e):** Export the highlighted events in the Events Flow View to a different file and load the new file in Events Flow View. 'e' is the keyboard shortcut for the action.
- **Save Currently Visible events only (v):** After zooming in to some extent, only the events in the visible area of the Events Flow View can be retained and loaded into a new file. 'v' is the keyboard shortcut for this action.
- **Text Mode (t):** Zoom in to see the text. In the text mode, the events are displayed with their attributes shown in lines. Only by zooming in to certain extent can you see the attribute lines. 't' is the keyboard short cut for this action. After you select **Text mode**, you can right-click in Events Flow View and select **Zoom to text** to see these attribute lines.
- **Show Diagnostic Events (d):** Show diagnostic events. 'd' is the keyboard shortcut for this action. After you selected this option, the event of debug interest is highlighted and crossed by a line, as shown in the following screen capture.




- **Gray Out Non-Highlighted Events (g):** Set the color of all the non-highlighted events to gray. 'g' is the keyboard shortcut for this action.
- **Highlight Thread Migration (m):** The multiprocessor view allows you to highlight and see the thread migrations. Too many unnecessary thread migrations from one processor to another can impact the performance in a negative way. 'm' is the keyboard shortcut for this action.

### Event Type Statistics view

The **Event Type Statistics** view shows the statistics for all or selected events by grouping them in the following hierarchy:

**Event Type>Event SubType>LoadModule>Module>offset**

The following screen capture shows the tree of the preceding hierarchy displayed in the **Event Type Statistics** view.



A tree node in the **Event Type Statistics** view has the following attributions:

- Event count percentage. This percentage is on the tree node and it means the percentage of the current event count to the total event count.
- Category name or offset value. According to different layers, the name can be the event type, event subtype, loadmodule, and module name.
- Event count.
- Event bars, which is shown in the following screen capture. The total number of the event bars equals the event count, and different colors of these bars denotes different loadmodules. You can get the event information by hovering over a bar.



The toolbar buttons of the **Event Type Statistics** view

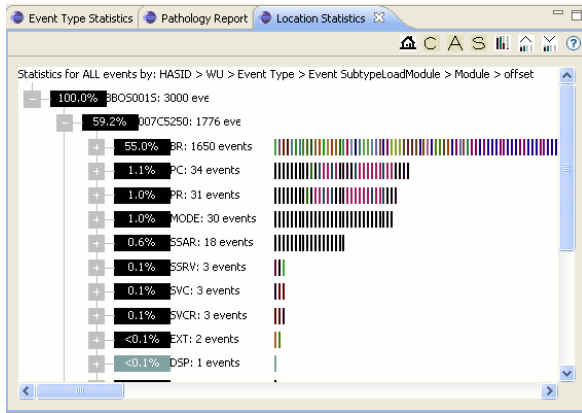
Name	Button	Description
	Home	Zoom to resize the tree to fit the window.
	Clear selected events	Clear all selections made in Events Flow View and all IBM MVS System Trace Analyzer views.
	Build statistics for all events	Build statistics for all the events in Events Flow View or other views.
	Build statistics for selected events	Build statistics for the selected events in Events Flow View or other views.
	Sort alphabetically	Sort the events in the alphabetically ascending order.
	Capture screen shot	Capture the screen shot of the current view.
	Show/hide elapsed time for paired events	Show or hide the elapsed time for paired events.
	Increase bar height for elapsed time	Increase the bar height that denotes the elapsed time of the matching events.
	Decrease bar height for elapsed time	Decrease the bar height that denotes the elapsed time of the matching events.

### Location Statistics view

The **Location Statistics** view shows the statistics for all or selected events by HASID and WU. The events are grouped in the following hierarchy:

**HASID>WU>Event Type>Event Subtype>LoadModule>Module>offset**

The following screen capture shows the tree of the preceding hierarchy displayed in the **Location Statistics** view.



A tree node in the **Location Statistics** view has the following attributions:

- Event count percentage. This percentage is on the tree node and it means the percentage of the current event count to the total event count.
- Category name or value. According to different layers, the name can be the event type, event subtype, loadmodule, and module name; the value can be the HASID, WU(work unit) and offset value.
- Event count.
- Event bars, which is shown in the following screen capture. The total number of the event bars equals the event count, and different colors of these bars denotes different loadmodules. You can get the event information by hovering over a bar.




The toolbar buttons of the **Event Type Statistics** view

Name	Button	Description
	Home	Zoom to resize the tree to fit the window.
	Clear selected events	Clear all selections made in Events Flow View and all IBM MVS System Trace Analyzer views.
	Build statistics for all events	Build statistics for all the events in Events Flow View or other views.
	Build statistics for selected events	Build statistics for the selected events in Events Flow View or other views.
	Sort alphabetically	Sort the events in the alphabetically ascending order.
	Capture screen shot	Capture the screen shot of the current view.
	Show/hide elapsed time for matching events	Show or hide the elapsed time for matching events.
	Increase bar height for elapsed time	Increase the bar height that denotes the elapsed time of the matching events.
	Decrease bar height for elapsed time	Decrease the bar height that denotes the elapsed time of the matching events.

## Pathology Report view

The **Pathology Report** view shows the the statistics for these events with abnormal activity.

The following screen capture shows all the events with abnormal behavior in the **Pathology Report** view, in which events are grouped by event type in a tree structure.



A tree node in the **Pathology Report** view has the following attributions:

- Event count percentage. This percentage is on the tree node and it means the percentage of the current event count to the total event count.
- Event type name, such as **EXT** shown in the preceding screen capture.
- Event count.
- Event bars, which is shown in the following screen capture. The total number of the event bars equals the event count, and different colors of these bars denotes different loadmodules. You can get the event information by hovering over a bar.



The toolbar buttons of the **Pathology Report** view

Name	Button	Description
	Home	Zoom to resize the tree to fit the window.
	Clear selected events	Clear all selections made in Events Flow View and all IBM MVS System Trace Analyzer views.
	Build statistics for all events	Build statistics for all the events in Events Flow View or other views.
	Build statistics for selected events	Build statistics for the selected events in Events Flow View or other views.
	Sort alphabetically	Sort the events in the alphabetically ascending order.
	Capture screen shot	Capture the screen shot of the current view.

## Sequence Context view

The **Sequence Context** view shows the event sequence of the trace file in the different types of context graphs, which are generated from the selected events in the Events Flow View or in the Event Type Statistics view. It answers the questions like: "How did I get to a certain type of events", or "What is the context of a hotspot". The view can display at most ten layers of event nodes (counted from the top highlighted node to the bottom leaf node of a graph).

To open the **Sequence Context** view, select **Window->Show View->Other**, and in the **Show View** dialog, select **IBM MVS System Trace Analyzer->Sequence Context**.

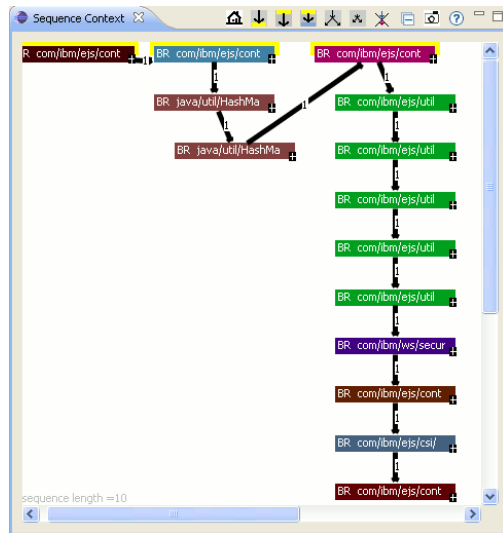
**Note:** The **Sequence Context** view works only when you installed Graphviz on your computer(see [Installation Graphviz](#) ).

The relationship between the parent and child nodes in both of the two graphs shows the order of the event occurrence. The event represented as a parent node always occur before the event represented as a child node.

There three types of context graph:


- In a forward context graph, the events occur in a sequence according to the path of the graph, starting from the root node, and you can also explore the child event nodes of the selected events furtherly.
- In a backward context graph, the events occur in a sequence according to the graph path, but you can trace back the parent nodes of the selected events in the reverse path of the graph.
- In a context graph from A to B, the events occur in a sequence according to the graph path, starting from the events highlighted in yellow to the events highlighted in blue in the Events Flow View or in other statistics views.

The following screen capture shows a forward context graph built from three selected events in the Events Flow View (see the second picture on this page for selected events). These selected events are built into graph nodes which are highlighted accordingly in the **Sequence Context** view, and the root event node starts from the selected event that occurs in the earliest time(in the range of your selection).



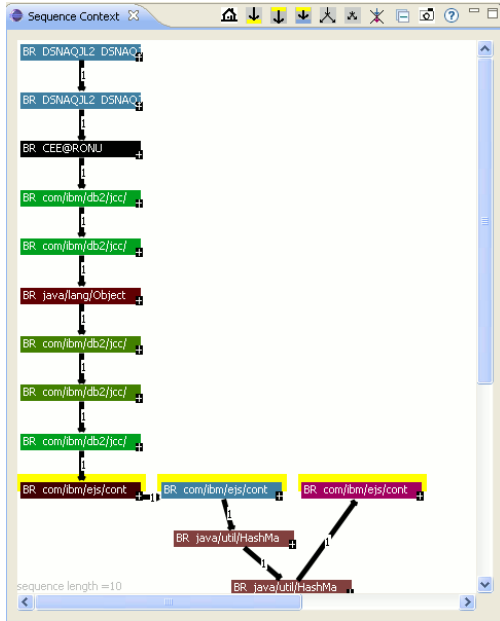
The text on each event node denotes the event name, and the number on the arrow denotes how many times the same events occur among the range of your selection in Events Flow View. For example, if the same event occurred twice among your range of selection, then the number is 2. Hover over the event node, and you can see the ToolTip, which shows the complete event name, and the number of the event occurrence.

The following screen capture shows the three selected events in the Events Flow View.




The following screen capture shows a backward context graph built from three selected events in the Events Flow View (see the second picture on this page for selected events). These selected events are built into the graph nodes which are highlighted accordingly in the **Sequence Context** view, and the leaf event node highlighted is the selected event that occurs in the latest time(in the range of your selection).





Each node in the diagram represents a group of events, each of which fits into the sequence drawn in the diagram.

You can inspect the paths between individual events by pressing Ctrl and clicking the little arrows. The following picture shows an example.



You can expand the node to show all the individual events by clicking the small + sign. Hovering over an event will reveal specific information about it. Of course, you can also click a node or an individual event to highlight the events across all views.



After expanding the individual events, you can inspect the path of an individual event by pressing Ctrl and clicking the little arrows. The following picture shows an example.



Some transitions between edges have an additional number between parentheses, which indicates the number of thread migrations. If you right-click this edge and select **Show all individual paths for this edge**, you can see some of the paths are drawn as a dashed line, meaning that they moved from one processor to another.

#### The toolbar buttons of Sequence Context view

Button	Name	Description
	Home	Zoom to resize the context graph to fit the window.
	Clear selected events	Clear all selections made in Events Flow View and all IBM MVS System Trace Analyzer views.
	How to get to the yellow events	Build backward sequence context graph according to the selection made in the Events Flow View or the statistics views.
	What comes after the yellow events	Build forward sequence context graph according to the selection made in the Events Flow View or the statistics views.
	How did we get from the yellow to the blue events	Build sequence context graph from the events highlighted in yellow to the events highlighted in blue.
	Make graph bigger	Zoom in.
	Make graph smaller	Zoom out.
	Clear the sequence paths between events	Clear the sequence paths between events.
	Collapse all nodes	Collapse all the expanded event nodes.
	Show/hide elapsed time for matching events	Show or hide the elapsed time for matching events.

	Increase bar height for elapsed time	Increase the bar height that denotes the elapsed time of the matching events.
	Decrease bar height for elapsed time	Decrease the bar height that denotes the elapsed time of the matching events.

### Statistics views

You can view the statistics of all or selected events with statistics views. The statistics views include the following ones:

- **Event Type Statistics** view
- **Location Statistics** view
- **Pathology Report** view
- **Sequence Context** view