

Defining and Using Machine Groups in the Performance Sentry PDB

An important feature of the Performance Sentry Performance Database (PDB) is the capability to categorize machines according to their function, usage, location or other salient factors and locate the performance data that is stored in the PDB associated with those machines. This document describe the basic facilities that are used to define Machine Groups and assign machines in the PDB to those groups, and also discusses best practices and procedures for using this capability effectively.

Assigning machines to Machine Groups an important function that is handled by the PDB Administrator. Machine groupings are an especially valuable feature whenever the PDB and its web reporting Portal are used to manage the performance data from hundreds or thousands of machines. In the Web Portal, Machine Groups assist in navigating within a large PDB, helping customers find machines of interest more quickly. Machine groupings also play a role in improving the performance of the reporting programs that the Web Portal provides, especially when the PDB needs to support a large number of machines.

Each machine available in the PDB can be assigned to as many as five Machine Groups, one Machine Group assignment per Machine Group category. There are five predefined Machine Group categories: *ApplicationGroup*, *LocationGroup*, *ClusterGroup*, *VirtualMachineGroup* and *WatchGroup*. You can define as many MachineGroups within each Machine Group category as are necessary to group the machines being monitored appropriately and help the end users of the PDB find machines of interest quickly.

There are several different ways machines are assigned to Machine Groups, including automatic assignment to an ApplicationGroup Machine Group when the first .smf data file from a machine is processed by the PDB Loader program, NtDACmd.exe. Machines can also be assigned to Machine Groups manually, using panels that are provided for that purpose in the Web Portal. Machine Group definitions and machine assignments to those Machine Group can also be performed in bulk using the NtDAUtil utility program.

When it is first installed, the PDB contains a small starter set of pre-defined Machine Groups, all in the ApplicationGroup category. These pre-defined ApplicationGroup Machine Groups are each associated with a Machine Group Mapping Definition that is also stored in the PDB. These mapping definitions are used to assign machines to these pre-defined ApplicationGroup Machine Groups automatically.

Machine Group Mapping Definitions associate one or more running Process instance names with an existing ApplicationGroup. When performance data for new machines is added to the PDB, these new machines are initially assigned to one of the pre-defined ApplicationGroup Machine Groups at runtime, by applying one of the mapping definitions. For example, when data from a machine is first loaded into the PDB, if the PDB Loader program discovers that a set of process level counters for an instance of the `sqlservr.exe` process exist, then that machine is automatically assigned to the pre-defined "SQL Server" ApplicationGroup. One of the first tasks of a PDB Administrator is to add to the starter set of pre-defined ApplicationGroup machine groups and the mapping definitions that assign machines to those machine groups automatically at run-time to better reflect the application mix that is running in your environment.

To use the Machine Groups feature of the NTSMF PDB effectively, you need to add to augment the predefined Machine Group definitions to reflect the specific Windows applications that your IT organization supports, the locations where processing is performed, the clustering of servers in those locations to support both high

availability and greater capacity, and the use of virtualization technology to allow server applications to make more effective use of underlying hardware resources. The additional Machine Groups that you define and assign machines to should reflect the priorities that are in place with the IT organization.

Only a few ApplicationGroup Machine Groups are pre-defined if you are installing a new PDB and only ApplicationGroup assignments are made automatically. Using Machine Groups effectively also requires adopting procedures to assign machines in the PDB to the Machine Groups you have defined. Operational procedures for maintaining the Machine Group assignments revolve around the following activities:

- Defining the Machine Groups that your installation will use to identify workloads and the machines that run those workloads.
- Actively assigning machines in the PDB to the Machine Groups you have defined.
- Modifying and extending the default set of Machine Group mapping definitions that automatically assign Machines to ApplicationGroup Machine Groups when performance counter data from a machine is first added to the PDB,

Defining the additional Machine Groups your organization will use and assigning machines to these Machine Groups are essential tasks for the Performance Sentry PDB Data Administrator. Note that the PDB Data Administrator panels that are used to define Machine Groups and extend the Machine Group mappings are only available in the Web Portal when the [EnableAdministrationPanels](#) Setting variable in web.config is set to a value of "Yes."

Basic Features

Each machine for which performance data is available in the PDB can be assigned to five separate Machine Groups, one per *Machine Group category*. The PDB provides support for five built-in Machine Group categories:

- [ApplicationGroup](#)
- [LocationGroup](#)
- [ClusterGroup](#)
- [VirtualMachineGroup](#)
- [WatchGroup](#)

Each Machine Group that you define must be associated with a Machine Group category, but there is no limit to the number of Machine Groups you can create in each of these categories. You can define as many Machine Groups in each category as are needed to support your organization and its IT priorities.

To get started, in the ApplicationGroup category, the PDB predefines several groups corresponding to Web servers, SQL servers, Exchange servers, etc. In addition, The PDB contains several pre-defined ApplicationGroup mappings that automatically assign a machine to an Application Machine Group the first time data for the machine is loaded into the PDB, based on the processes that were found to be running on the machine at the time the initial is loaded. The pre-defined ApplicationGroup mappings are shown below in Table 1, alongside the name of the process that is used to assign a machine automatically to that ApplicationGroup.

Table 1. Pre-defined ApplicationGroup process mappings.

ApplicationGroup Machine Group	Maps to process...
SQL Server	sqlservr
Web Server	w3wp

Exchange Server	store
Domain Controller	ismserve
Oracle	oracle

Note that you are not limited in automatically assigning a machine to an Application Group machine group to a single identifying process instance name. You can define two or more process names that all map to the same ApplicationGroup domain. You can define additional ApplicationGroup Machine Groups as necessary, and define additional mappings to have these Application Machine Groups assigned automatically to machines when counter data from these machines is first loaded into the PDB.

You can use the other Machine Group categories to group machines in a variety of useful ways. Within the WatchGroup category, for example, you could create Watch groups for each performance analyst (Bill's Hot list, Sarah's Priority Watch, etc.) or based upon current performance investigations. The only limitation is that, within each of the five Machine Group categories, a given machine can only be assigned to one Machine Group at a time.

Utility functions are available in the Web Portal to define Machine Groups and assign machines to those groups. However, instead of relying on a manual process to assign large numbers of machines to groups, you can automate the process using the NtDaUtil.exe program's **-loadgroups** option, which loads both the Machine Groups and Machine Group assignments from an .xml format file. Additional Machine Group utilities in the web Portal are available that are designed to support this automation procedure. These additional facilities to assist in Machine Group assignment automation include:

- an Export utility to write the current Machine Group definitions and machine group assignments to an xml format file
- an Unassigned Machines report to identify machines without any current Machine Group assignments

DefaultMachineGroup setting in the Web Portal

The Web Portal requires the use of Machine Group definitions in key spots to speed up reporting programs. For example, in building the Alerts daily summary report, the online report is limited to scanning the PDB for alerts only for machines that belong to a specific Machine Group. To create the Alerts daily summary report, the PDB's Alert events table and the Alert config table for each machine to be included in the report is queried. Since scanning the entire PDB can be expensive, the Portal limits the scan to a set of machines in a designated Machine Group.

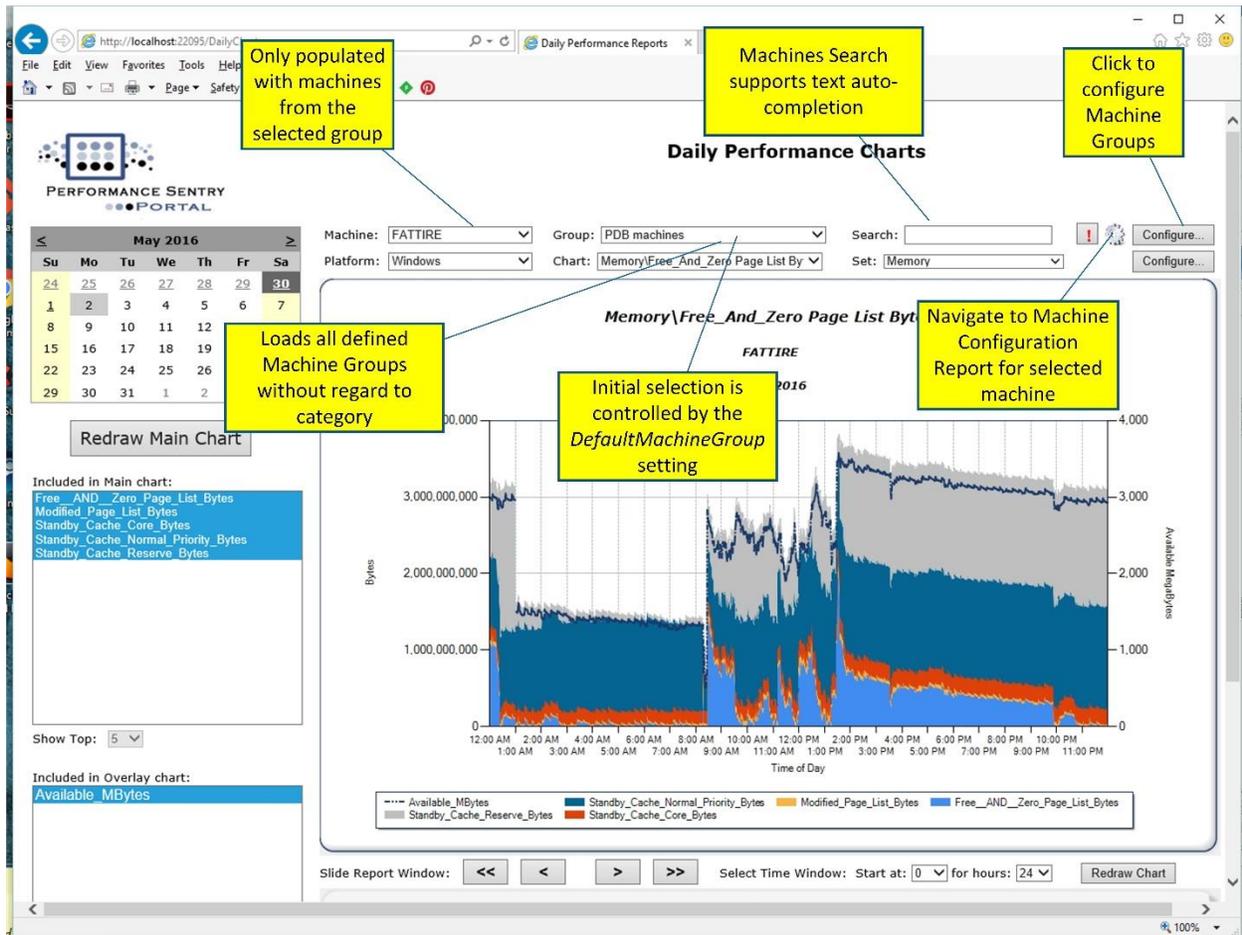


Figure 1. Screen shot of the upper portion of the Daily Performance Charts web page, with annotations showing how Machine Group definitions and the machines assigned to those groupings are integrated into the panel.

The screen shot of the Daily Performance Charts web page shown in Figure 1 illustrates several ways that Machine Groups are integrated into the Web Portal UI. Machine Group definitions are used to construct the machine selection menus that are used in the Daily Performance Charts (as illustrated in Figure 1), helping that web page to load faster. Building a Machines selection menu on the web server that encompasses every machine defined to a large PDB and sending to the web browser can be prohibitively expensive, contributing to slow initial Page Load times.

Note that on the Daily Performance Charts web page, the Machine Group dropdown selection menu loads all defined Machine Groups without regard to the Machine Group category. The machine and machine group selection and search controls on the Server Health reporting page and the PDB Data Administration landing page

(illustrated by the screen shot shown in Figure 2) conform to the same pattern as the Daily Performance Charts web page.

The Portal uses an *application setting* to identify a default Machine Group, corresponding to the Machine Group that is marked as initially selected in the machine selection menus. This app setting is used by several of the web pages, including the Daily Performance Charts web page illustrated in Figure 1, when they first load to select the initial set of machines that populate the page's machine selection controls.

The default Machine Group can be *any* MachineGroup that you have defined. To identify the default Machine Group, you can edit the **DefaultMachineGroup** setting in the <appSettings> section of the Portal's **web.config** file by hand. Alternatively, you can use controls in the Web Portal in the PDB Database Administration pages to set the default Machine Group.

Alerts Daily Summary report. If the Machines table contains more entries than the **MaxMachinesToScanForAlertReporting** setting, the Alert Daily Summary reporting program will locate and process alerts from the *machinename@AlertEvent* table initially only for machines in the default Machine Group. The Alert Daily Summary reporting page, Alerts.aspx, also contains a Machine Group selection menu that allows you to generate additional Alert Daily Summary reports on a machine group by machine group basis.

Machine selection controls. Machine Groups are used to speed up the initial time it takes to load both the Daily Performance Report, DailyCharts.aspx, the Server Health report, and the PDB Data Administration landing page, PDBDataAdministration.aspx. All these web pages feature machine selection menus that are populated initially from the default Machine Group. On each of these pages, these machine selection menus are augmented by a Machine Group dropdown and a Search capability that incorporates textbox auto-completion hints, computed based on the first few characters that are typed into the Search box.

The **DefaultMachineGroup** setting determines which of the Machine Groups that you have defined is initially selected in the Machine Group dropdown menu on these web pages. The default Machine Group is used to populate an initial Machines selection dropdown on these web pages. The Web Portal programs query the PDB to enumerate the machines found in the PDB that are assigned to the selected Machine Group. The result set from this PDB query is then used to populate the Machines selection dropdown. The first machine found in the default Machine Group is then selected by default. Following the initial Page Load, the Machine Groups and Machines dropdown selection menus are used to navigate to a specific machine.

Recommended Practices and Procedures.

When you are first installing the Performance Sentry PDB and begin populating it with performance counter data from your Windows, VMware and Linux machines, it should not be necessary to expend a great deal of effort to define Machine Groups or follow strict procedures to ensure that all machines are assigned to one or more machine groups. This is especially true when you are in the Proof of Concept phase where you are still evaluating the product. However, for any machines that you would like to reference quickly when you are accessing Daily Performance Charts, for example, as Figure 1 illustrates, they need to be assigned to at least one Machine Group. This is, of course, the reason that an Application Machine Group is assigned automatically to each machine initially during a PDB load.

With a relatively small PDB, you can define and assign machines to other Machine Groups manually on an *ad hoc* basis, using the PDB Data Administration screens that were designed for this purpose. However, over time as you begin to rely on the PDB for storage and retrieval of your performance counter data and the number of machines you are managing with the Performance Sentry PDB starts to grow, using Machine Groups effectively is a key

factor in using the NTSMF PDB successfully in your organization. At this point, it is important to determine what Machine Groups you will need in each of the five pre-defined Machine Group categories. Remember there is no limit to the number of Machine Groups you can define. If your IT organization has strong server naming conventions in place, for example, the machine Search function may be so easy to use that PDB customers may not need an elaborate set of Machine Group definitions and machine assignments to navigate to machines of interest. When that is not the case, however, defining a strong set of Machine Groups is necessary to help Portal users navigate within a large PDB that contains data from many different machines.

This section discusses our recommended practices and procedures for configuring Machine Groups and keeping Machine Group assignments up-to-date.

Getting started.

The PDB Database Administration section of the web Portal contains pages you can use online to define Machine Groups and assign machines to those groups. From the Daily Performance Charts page, click the “Configure” button across from the machine selection controls to access the main PDB Database Administration page. See Figure 2 for an example of the PDB Database Administration page. The buttons on the left side of the panel, underneath the heading “Manage Machine Settings,” are used to define Machine Groups and assign machines to those groups.

The first set of actions you should take is to define the additional Machine Groups in the Watch, Application, Location, Cluster, and Virtual Machine Group categories that will help you organize the machines in the organization that you will be reporting on. Click the button labeled “Define Machine Groups” to view the current set of Machine Group definitions and to define new ones. When defining a new Machine Group, you must always associate it with one of the five Machine Group categories.

PDB Database Administration

Daily Operations Summary Report

	Date	Total Files	Errors	Thruput (Records/sec)	Avg. Time Per File (secs)	Avg. File Size (records)
View Log	1/24/2014	2	0	550	302	166017
View Log	1/26/2014	1	0	240	1697	407272
View Log	1/27/2014	2	0	536	332	177897
View Log	1/28/2014	1	0	361	1207	436010

Manage Machine Settings

Machine:

Group:

or

Search:

No file chosen

File Name:

Figure 2. The PDB Database Administration landing page. The machine selection controls under the “Manage Machine Settings” heading are used to define Machine Groups and assign machines to those groups.

If you click the “Define Machine Groups” button on the PDB Administrator’s landing page shown in Figure 2, you will navigate to the panel where you can define additional Machine Groups and ApplicationGroup Machine Group mappings. An example of the web page for defining Machine Groups and Application Machine Groups is shown in Figure 3.

Figure 3. The PDB Database Administration page for defining Machine Groups and ApplicationGroup Machine Group Mappings.

Define New Machine Groups:

New Group Name: Group Type: WatchGroup

Define Application Machine Group mappings:

Group name: Hyper-V Servers vmwp

Machine selection (for Map to Process query)

Machine: WS2012VMHOST

explorer
mmc
mspaint
NTDACmd
svchost
System
vmconnect
vmwp

All Machine Groups:

Group Name	Group Type	RetainDetailDays
Domain Controller	ApplicationGroup	
Exchange Server	ApplicationGroup	
Citrix Server	ApplicationGroup	0
Hyper-V Servers	ApplicationGroup	0
Middleware	ApplicationGroup	0
Sentry VM	ApplicationGroup	0
Sharepoint Servers	ApplicationGroup	0
SQL Server	ApplicationGroup	
Web Browser Application Group	ApplicationGroup	0
Web Server	ApplicationGroup	
Workstations	ApplicationGroup	0
SW Florida	LocationGroup	0
Naples FL	LocationGroup	0
AWS	LocationGroup	0
Puget Sound Area	LocationGroup	0
VMware ESX Hosts	VirtualMachineGroup	0

12

Default Machine Group: PDB machines

The top section of the Define MachineGroups page has a set of controls for defining new Machine Groups, each of which must be associated with one of the five Machine Group categories.

ApplicationGroup Machine Group mappings

The middle section of the web page is used to define new ApplicationGroup Machine Group mappings. The PDB contains an initial set of ApplicationGroup Machine Group mappings that are used to assign a machine to an ApplicationGroup Machine Group automatically when data from the machine is first loaded into the PDB. The initial, pre-defined ApplicationGroup Machine Group mappings were shown earlier in Table 1. When the automatic assignment function executed during an initial PDB load cannot find any of the processes that map a machine to an ApplicationGroup, the machine is assigned to a catch-all “Other” category. Whenever too many of your machines remain assigned to this undifferentiated “Other” category, you probably need to define additional ApplicationGroup Machine Group mappings.

You can use the web page controls shown here in Figure 3 to define additional ApplicationGroup Machine Group mappings. The Group name dropdown displays each of the current Application Machine Groups. You can only define new ApplicationGroup Machine Group mappings for ApplicationGroup Machine Groups that are currently defined. First, make sure the ApplicationGroup Machine Group that you want to add process mappings for currently exists.

For each Application Machine Group, you can define one or more process instance mappings. These mappings control the initial assignment of a new machine to one of the Application Machine Groups. Note that you can always manually change the ApplicationGroup Machine Group that was assigned automatically to the machine when counter data from the machine was first processed and added to the PDB. Automatic assignment of an ApplicationGroup Machine Group only occurs when the ApplicationGroup field in the Machines table entry for the machine is blank. If you need to override the automatic assignment of a machine to an ApplicationGroup Machine Group, see the instructions for [“Assigning Machines to Machine Groups”](#) in the following section.

Additional controls are provided on the web page to make defining new mapping definitions less error-prone. They allow you to select a machine currently defined to a designated Application Machine Group and execute a PDB query that returns a list of process instance names from the selected machine. Once it is populated with process instance names from the PDB, when you click on a process name in the process instance List Box, the process name you select is copied to the “maps to process” input field. Then when you press the “Add Mapping” button, the Application Machine Group mapping you have defined is added to the PDB.

The lower portion of the page contains a tabular display of each of the current Machine Groups that are defined.

At the bottom of the page is a set of controls for setting the Default Machine Group. The Default Machine Group is used initially in the Daily Charts, Server Health Reports, the Alerts Reports, and elsewhere in the PDB Administration screens. For more information on the Web Portal’s Default Machine Group, see the section of this document entitled [“DefaultMachineGroup setting in the Web Portal.”](#)

Assigning Machines to Machine Groups

After you have defined one or more Machine Groups, you can start assigning machines to them. Only machines that already have counter data loaded in the PDB can be assigned to Machine Groups. Loading data from a machine in the PDB for the first time creates an entry in the Machines table for that machine, and the Machine Group assignments for a machine are stored in each machine’s Machines table entry.

Using panels in the Web Portal, you can assign machines to Machine Groups in several ways. If you only want to assign one specific machine to one or more Machine Groups, access the Machine Configuration report. To begin, select a specific machine using the machine selection controls under the “Manage Machine Settings” heading on the main PDB Data Administration page (illustrated in Figure 2). Select a machine by name from the Machine

dropdown or use the Search control to find a specific machine and then press the big button that displays a magnifying glass. Clicking this button transfers control to a Machine Configuration report page where you will see controls that describe the current Machine Group settings for that machine and an Assign button that allows you to assign new or modify any of the five existing Machine Group settings. Using the controls on the Machine Configuration report, you can assign a machine to a Machine Group in each of the five Machine Group categories by clicking the large “Assign Machine Groups” button.

Another way to assign machines to Machine Groups is to click the “Assign Machines to Groups” button on the PDB Data Administration landing page shown in Figure 2, which transfers control to a page where you can assign multiples machines to a Machine Group, one Machine Group at a time. The Machine Group Assignments page, MachineGroupAssignments.aspx, builds a report that identifies each machine in the Machines table and indicates the current Machine Group assignment in all five Machine Group categories. The tabular report that is built can be sorted by clicking on any of the column headers. You can also filter the list of machines that are reported, as illustrated in Figure 4, select these machines, and then assign all the selected machines to a specific Machine Group by clicking the “Assign” button.

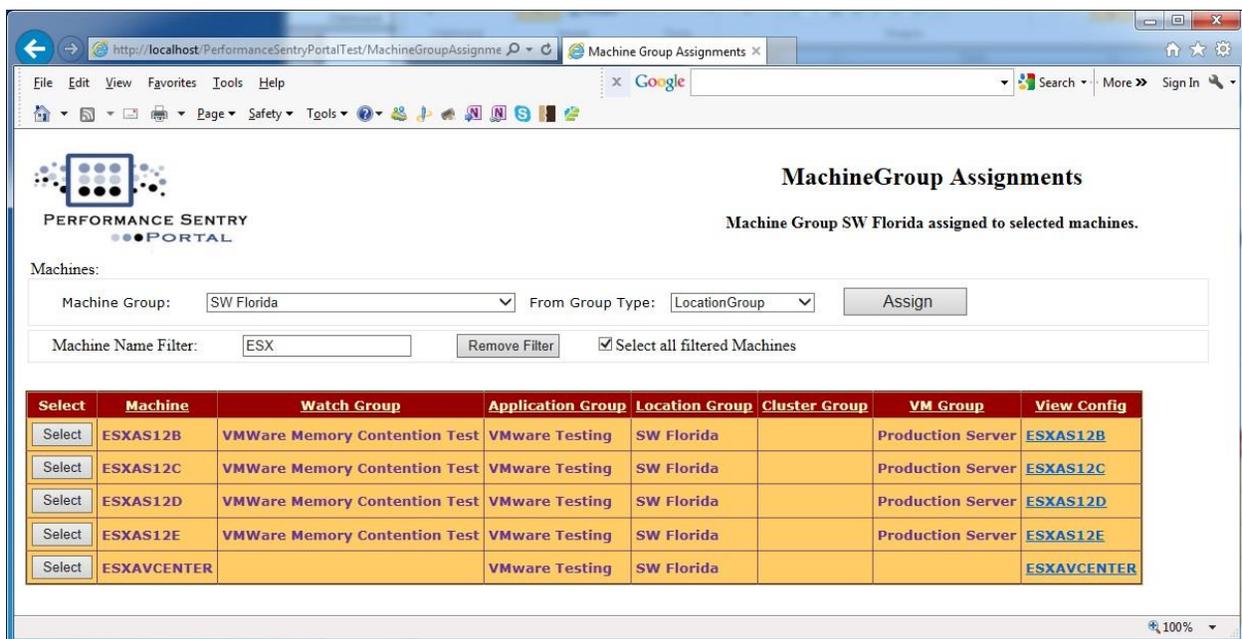


Figure 4. Machine Group Assignments page in the web portal, illustrating specifying a Machine name filter, then clicking the “Assign” button to assign a group of related machines to the same Machine Group.

Finally, you can create a Machine Group definition file in xml format and run the NtDAUtil **-loadgroups** utility to both define Machine Groups and assign machines to those groups. This is usually the preferred option for PDB DBAs that are responsible for a large configuration of machines. The procedures for automating the Machine Group assignments using a Machine Group definition xml file are discussed in the next section.

Machine Group assignment automation.

Once the Performance Sentry PDB is up and running and you are pushing counter data from hundreds or thousands of machines into the database on a daily basis, you will want to automate the process for maintaining the Machine Group assignments.

To get started, generate a Machine Group definition file in xml format, reflecting the current configuration. Click the “Export Machine Group settings” button at bottom left on the PDB Database Administration landing page to

create a Machine Group definition file in xml format. You can then edit and make changes to this Machine Group definition file.

Then, you can run the NtDAUtil **-loadgroups** utility. Note that the NtDAUtil **-loadgroups** utility will *replace* all current Machine Groups definitions and machine assignments using the contents of the specified Machine Group definition file, so make sure you have a current back-up of the existing configuration.

The Machine Groups definition xml file is in the following format:

```
<GroupNames>
  <Group GroupName= PriorityWatch >
    <GroupType>WatchGroup</GroupType>
    <RetainDetailDays>120</RetainDetailDays>
    <Description>This is my high priority WatchGroup</Description>
    <Machines>
      <Machine MachineName= 27NN943L />
      <Machine MachineName= TS-CATV5-220 />
      ...
    </Machines>
  </Group>
  <Group GroupName= HQ Prod>
    <GroupType>LocationGroup</GroupType>
    <RetainDetailDays>60</RetainDetailDays>
    <Description>This is my HQ Production LocationGroup</Description>
    <Machines>
      <Machine MachineName= 27NN943L />
    </Machines>
  </Group>
  <Group GroupName= Test Servers HQ>
    <GroupType>LocationGroup</GroupType>
    <RetainDetailDays>60</RetainDetailDays>
    <Description>This is for Test Servers located in HQ</Description>
    <Machines>
      <Machine MachineName= TS-CATV5-220 />
    </Machines>
  </Group>
  ...
</GroupNames>
```

For each Group defined in the Machine Groups definition xml file, you must specify its **GroupName** property and **<GroupType>** attribute, as illustrated. The **<RetainDetailDays>** attribute allows you to set a database retention policy at the Machine Group level that overrides the system-wide default, which is defined in the Registry at *HKLM\SOFTWARE\DemandTechnology\PerformanceSeNtry\DataAccess\RetainDetailDays*. Optionally, the definition file can also enumerate the machines that are assigned to each MachineGroup. When you want to assign a specific machine to more than one Machine Group, just add a **<Machine>** entry to the list of **<Machines>** that are defined for each Machine Group as appropriate.

In building the Machine Groups definition xml file, it is recommended that you use Export the utility on the Web Portal's main PDB Database Administration landing page to get started based on the current configuration.

Unassigned Machines

To keep your machine group assignments current, access the Unassigned Machines report from the main PDB Database Administration landing page. This report lists any machines that are not assigned to any Machine Groups, and it provides controls to assign those machines to groups, one machine at a time. As illustrated in Figure 5, set the dropdown lists for each Machine Group category and then click the “Modify” in the associated row next to the Machine name. Click on the machine name in the “View” column at the far right of the table to generate the Configuration and Inventory report for that machine, where, as discussed above, you can also assign the machine to Machine Groups. Accessing the Machine Configuration report page is sometimes helpful when you are not familiar with the machine or what applications it is running.

An example of the Unassigned Machines report is shown in Figure 5. Periodically, you should run the Unassigned Machines report to see if the PDB contains data from any machines that are not currently assigned to any Machine Groups, making them difficult to access for reporting purposes. If you make any changes to the Machine Group assignments online using the Unassigned Machines report, be sure to re-run the Export utility to update your offline copy of the Machine Group definition xml file with the latest Machine Group assignments.

The screenshot shows the Performance Sentry Portal interface. At the top, there is a navigation bar with 'File', 'Edit', 'View', 'Favorites', 'Tools', and 'Help'. Below this is a search bar and a 'Sign In' button. The main content area is titled 'Unassigned Machines Report'. There is a 'Define Machine Groups' button and a 'Default Machine Group' dropdown menu set to 'Priority Watch'. Below this is an 'Assign' button. The main table has the following columns: Machine, Watch Group, Application Group, Location Group, Cluster Group, VM Group, and View. The table contains 15 rows of machine data. The 'Application Group' for the machine 'SMF-150-0-15' is highlighted in blue and set to 'NTSMFPDB machines'.

Machine	Watch Group	Application Group	Location Group	Cluster Group	VM Group	View
APP-STL-1351	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	APP-STL-1351
BCS-PPDGTW-03B	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	BCS-PPDGTW-03B
DELLT3400X64	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	DELLT3400X64
GN3F843Z	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	GN3F843Z
S3W04619	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	S3W04619
S3W06161	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	S3W06161
SE105605	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	SE105605
SE107425	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	SE107425
SE107426	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	SE107426
SMF-150-0-15	Priority Watch	NTSMFPDB machines	Customer samples	* Unassigned *	* Unassigned *	SMF-150-0-15
SMF-150-0-16	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	SMF-150-0-16
SXLOE104	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	SXLOE104
WIN-H7109KPS0RL	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	WIN-H7109KPS0RL
WS08R2X64	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	* Unassigned *	WS08R2X64

Figure 5. The Unassigned Machines report. Periodically, you should run the Unassigned Machines report to see if the PDB contains data from machines that are not currently assigned to any Machine Groups. Whenever you make changes to the Machine Group assignments online using the Unassigned Machines report, you should then re-run the Export utility to refresh your offline copy of the Machine Group definition xml file.