



- The **zManager** provides basic performance functions and performance data collection for all workloads for which you have activated service policies.
 - With additional customization, you can benefit from more advanced performance management functions and collect *more granular performance data* to help identify and resolve hybrid performance problems.
 - These advanced capabilities are provided through the use of guest platform management providers that you install or start on virtual servers in the ensemble.
- After creating a workload and activating a performance policy for it, zEnterprise performance management functions provide the following basic support:
 - Monitoring and data collection at the **hypervisor** and **virtual server levels**, with data provided through these reports: Service class report, Virtual server report, Hypervisor report, Resource adjustment reports by workload, virtual server and service class reports.
 - Performance management according to the velocity goals that you defined for the service classes in the active policy.
 - CPU management that enables resource allocations to be adjusted dynamically between virtual servers that have the 'same' hypervisor.

- Additional customization includes installing or starting **guest platform management providers (GPMPs)** where you can benefit from advanced capabilities and additional performance data:
 - Virtual server topology report, including operating system statistics (a)
 - Hops report for a service class (b)
- You can collect application data from IBM middleware products that have been instrumented with the Open Group **Application Response Measurement (ARM) 4.0** standard.
 - [Guest platform management providers are supported on the following virtual servers] :
 - The logical partitions (LPARs) running z/OS Version 1 Release 12 (V1R12), V1R11, or V1R10.
 - The z/VM guests running z/OS Version 1 Release 12 (V1R12), V1R11, or V1R10.
 - The z/VM guests running Linux.
 - The virtual servers running AIX on a POWER blade.

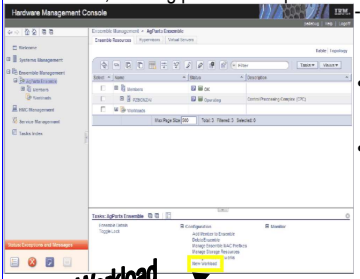



- Advantages of configuring application response measurement (ARM)**
- To ensure that work requests are performing as expected in a **multi-tiered, heterogeneous server environment**, you can **track the performance** of those requests across server and subsystem boundaries, and manage the underlying physical and network resources used to achieve specified performance goals.
- Note:** ARM standard is an application programming interface (API) that allows you to monitor the availability and performance of business transactions within and across diverse applications and systems.
- Within a zEnterprise ensemble, you can collect the following additional performance data from IBM ARM instrumented middleware:
 - The **response time** of work requests in the instrumented applications
 - An **end-to-end** view of work requests as they flow through the ensemble
 - The **topology** of the systems that perform the work requests for the instrumented applications.



NOTE: If you plan to use the guest platform management provider, there are actions you should take at implementation time. The **gpmp** command displays or modifies the guest platform management provider, which collects performance data for ARM-instrumented applications or processes. ARM-instrumented applications or processes use application programming interfaces (APIs) according to the Open Group Application Response Measurement (ARM) standard.

- This infrastructure includes the following software running on a virtual server:
 - ARM-instrumented middleware applications**, which include **DB2** and **WebSphere Application Server** products that run on **z/OS, AIX or Linux**. These middleware applications use ARM API calls while servicing requests from multiple clients.
 - The **ARM 4.0 Library Routine**, which exports the full set of ARM 4.0 interfaces as defined by the Open Group standard. Instrumented applications dynamically bind to the ARM library routine that is distributed with a particular operating sys.
 - The guest platform management provider (GPMP), which collects performance data for work running on a virtual server. The GPMP *passes the data* to the ensemble HMC, where you can view the data in various reports.
- The **New Workload Wizard** is an HMC guide that steps you through the process of defining a workload, creating performance policies and service classes, and activating one policy.




New Workload

- The **New Workload Wizard** is particularly useful when you are first setting up an ensemble and its elements, including policy definitions.
 - It also provides a consolidated overview of the workload-related tasks that you can perform separately, during normal ensemble operations.
 - As an example, for **pBlades** after you have determined the application resource requirements, you start the New Virtual Server wizard and go through the panels to define the new server characteristics.
 - Two panels** in the New Virtual Server wizard are **key** to performance management:
 - On the **Assign Processors** panel, you need to specify whether your new virtual server is to share or have dedicated use of the blade processors.
 - Note:** When you select Shared as the processing mode, you also specify the number of blade resources, called *virtual processors*, that a new virtual server uses.
 - On the **Performance Management** panel, you select Enabled to allow the zManager to manage CPU resources for this new virtual server.

- To determine whether you need to adjust the default values that are assigned to a virtual server on a **POWER blade**, you need to consider the following factors.
 - The **initial processing capacity** required by the applications that are to run on this virtual server.
 - The effect of **running the virtual server** at the minimum rather than the initial capacity value.
 - The virtual server's **role as donor or receiver** when processor management is enabled.
- Note:** For **z/VM**, the zManager and the z/VM Resource Manager (VMRM) can manage CPU resources among z/VM virtual servers, but only one can perform this function at a time.

Creating and managing workloads

- Administration/operation tasks accomplished through ensemble-management HMC.
 - Create a new workload context
 - Edit workload context
 - List workload contexts and associated policies
 - Delete workload context
 - View virtual server property
- When you enable CPU management in an ensemble, the IBM zEnterprise Unified Resource Manager (zManager) uses goal-oriented policies to manage CPU resources that are allotted to virtual servers in the ensemble.
 - Resources can be adjusted only among virtual servers that run under the *same* hypervisor.
- CPU management is supported for z/VM hypervisors running on the z196, and for PowerVM Enterprise Edition hypervisors running on a POWER blade.
 - Before the zManager can dynamically adjust CPU resources according to goal-based policies, **CPU management must be enabled** for both the hypervisor type and for individual virtual servers.
 - When you first create an ensemble, the initial setting for CPU management for hypervisor types is **Disabled**.
 - If you decide to enable CPU management in the ensemble for a hypervisor type, you need to make some configuration changes to the z/VM or PowerVM hypervisor before doing so.
 - When you first **create a virtual server**, the initial setting for CPU management is **Enabled**; however, the zManager does not manage CPU allocations unless CPU management *also* is enabled for that hypervisor.



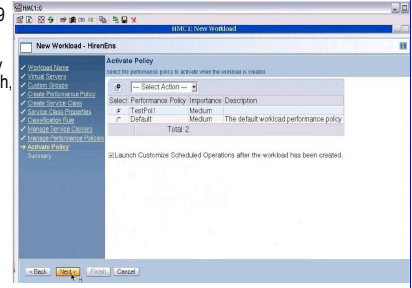
Prerequisites for using performance management

- The performance management functions that are available for use in a zEnterprise environment depend on the suite associated with the zManager (**MANAGE** or **AUTOMATE**).
 - The basic **Manage** suite provides **limited** performance management functions.
 - All virtual servers in the zEnterprise and IBM zEnterprise BladeCenter Extension (zBX) are automatically associated with the **default workload** and performance policy.
- Through the Monitors Dashboard in the HMC, you can display CPU and I/O activity.
 - The **Automate** suite provides the high-value management functions that enable you to create your own custom workload and performance policies.
 - With the **Automate** suite, you are able to manage workloads by:
 - Defining a workload and associating specific virtual servers with that custom workload.
 - Creating performance policies for each custom workload.
 - Monitoring and displaying performance data and other details for each workload.
 - Monitoring platform resources used to support each workload.
 - Viewing overall workload performance health from a platform perspective.
 - Displaying whether goals defined in the workload performance policy are being achieved.
 - Drilling down to identify which virtual servers are contributing to performance problems.

Note: Also with the Automate suite, your custom workloads can benefit from dynamic adjustments to CPU resources to ensure that multitenant applications are provided sufficient resources, and adhere to service level agreements. **Note:** Performance management functions provided through either the Manage or Automate suite **do not** apply for the IBM Smart Analytics Optimizer for DB2 for z/OS.

Elements of a workload management policy

- Importance (Highest, High, Medium, Low, Lowest)
- Describe the connection between Workload level importance, and service class level importance.
 - Service class
 - More than one service class can be defined for a policy; max 99 service classes per policy
 - Performance goal
 - Velocity (Fastest, Fast, Moderate, Slow, Slowest) ;Discretionary
 - Business importance (Specified with Velocity only: Highest, High, Medium, Low, Lowest)
 - Classification rule
 - More than one rule can be defined for a service class
 - Filter attribute
 - > Hostname
 - > OS Level
 - > OS Name
 - > OS Type
 - > Virtual Server Name
 - Filter operator
 - Filter value
 - Values for default performance policy
 - More than one policy can be defined for a workload, but only one can be active at a time.
 - Export / import policy functions (to-from ISPF WLM)



There is no concept of Service Class periods. It's all under one period similar to long running tasks.

To collect data from ARM-instrumented applications, the IBM Systems Director Enterprise Workload Manager (EWLM) must be enabled.