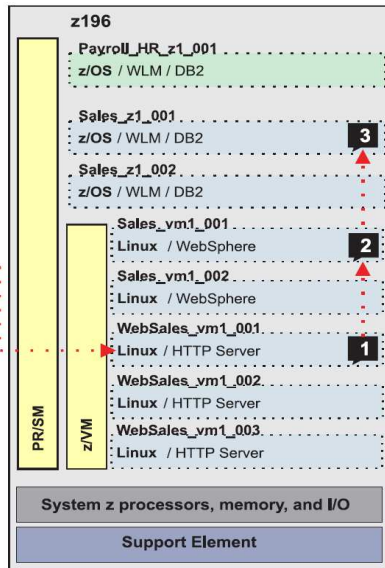


Cheat Sheet #76 zTidBits zEnterprise Cooperative Management and zOS WLM

The performance management functions that are available for use in a zEnterprise runtime depend on the suite associated with the Unified Resource Manager.

Figure 1




When requirements are met, the Unified Resource Manager workload performance policy service class is mapped to a z/OS WLM service class when work enters the z/OS runtime.

From that point, the work on the z/OS tier benefits from the full range of z/OS WLM management capabilities.

To monitor virtual server performance, use the monitoring reports that are available through the primary HMC; to monitor work segment performance on z/OS, use z/OS monitoring tools such as Resource Measurement Facility (RMF™).

See: #27 zNibbler (zEnterprise Platform Performance Management)

Figure 2



- Through a workload performance policy, you can define goals that the Unified Resource Manager uses to manage virtual servers.
- For virtual servers that run z/OS, you can apply these performance goals to the work requests that z/OS and its subsystems support.
- When the z/OS environment includes properly configured guest platform management providers and instrumented IBM middleware products, you can map ensemble workload performance goals to z/OS workload manager (WLM) service classes, thus achieving end-to-end goal-based performance management for multitier applications.

Classification of work requests in a non-zEnterprise environment

For example, suppose that you use a composite application to manage Internet sales. This application consists of web-services-enabled vendor products and IBM middleware in a three-tier structure shown in figure 1 below - right. These vendor products and middleware are installed on operating systems that run on a z196 that is not an ensemble member.

- An instance of IBM HTTP Server (tier 1) receives the incoming work request and routes it to an instance of WebSphere® Application Server (tier 2). The instances of HTTP Server and WebSphere Application Server run on z/VM Linux guests.
- WebSphere Application Server routes the work request to an instance of DB2 running on a z/OS partition (tier 3).
- On z/OS, the incoming work request is classified as a DB2 transaction and z/OS WLM manages it according to the service class for that type of work request.

- In a non-zEnterprise environment, your company has to use a collection of separate products or components such as z/OS WLM to manage work requests within each tier.
- Each product or component has no knowledge of the performance goals that the other products or components are using to manage the work segments.
- When work enters z/OS, WLM compares information about each work request to classification rules that a z/OS WLM administrator has defined.
- Classification rules are filters that you define to categorize work into service classes, and optionally report classes, based on work qualifiers. A work qualifier is information that identifies a work request to the system.

For this example, the z/OS WLM administrator has defined two service classes for DB2 work requests: DB2TRANS for most incoming work and DB2BANK for those transactions with the name CREDIT. These two service classes have different response time goals set for managing the work requests. Figure 2 provides a conceptual illustration of the classification process.

- The box on the left side of the figure shows partial information that is available when the work request enters z/OS. WLM correlates this information with the work qualifiers in the defined classification rules.
- These qualifiers identify the DB2 subsystem instance (SI) and the transaction name (TN) to determine that DB2BANK, rather than DB2TRANS, is the appropriate service class for this work request. At this point in the example, WLM classifies the work request without any knowledge of the workload performance goals that are defined for the virtual servers that support the multi-tiered application.

Classification of work requests in a zEnterprise ensemble

- Because the virtual servers that support this multitier application are assigned to the same workload, the work request is associated with the ensemble workload policy that the Unified Resource Manager uses to manage these virtual servers.
- Figure 3 shows that the virtual servers belong to the Product Sales workload.
- The active performance policy is Peak, which has a WebSales service class to which the virtual servers are defined through classification rules based on the virtual server name (hostname).

For example, the classification rule Hostname==Sales* identifies the two z/OS virtual servers and two of the Linux guests running on z/VM. For the work that these virtual servers support, the WebSales service class sets the business importance as high and the velocity goal as fast.

- To apply these the Unified Resource Manager workload policy goals to work segments in the z/OS tier, z/OS WLM administrators can define new WLM service classes and classification rules specifically for work requests that the z/OS virtual servers support.

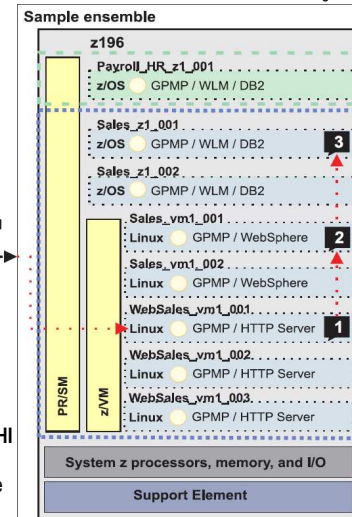


Figure 3

For this example, suppose that the z/OS WLM administrator defines two new service classes: ZMGRDEFA is the default service class for most work and GPFASHTI is for those transactions with the Unified Resource Manager service class name of WebSales.

- Through these two WLM service classes, the z/OS WLM administrator defines different response time goals for the incoming work requests arriving from other virtual servers that are in the same ensemble as the z/OS virtual servers.
- Figure 4 provides a conceptual illustration of the new classification process.
- When the work request enters z/OS, WLM correlates the information about the work request with the work qualifiers in the defined classification rules.
- The ESC work qualifier identifies work requests as associated with a zManager workload performance policy and with the service class to which the z/OS virtual server is defined.
- In this example, the Unified Resource Manager service class name WebSales identifies GPFASHTI as the appropriate service class for this work request. In effect, this classification process maps the ensemble workload performance service classes to WLM service classes, which enables end-to-end performance management for the work requests that z/OS virtual servers support.

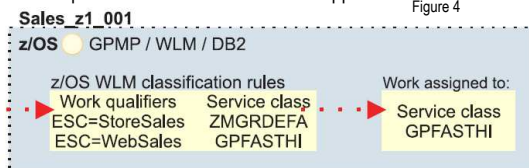


Figure 4

Mapping Unified Resource Manager (URM) service classes to WLM services classes requires these tasks:

For the ensemble

- Create a workload and add the z/OS virtual servers to it. Also enable guest platform management provider (GPMP) support for each z/OS virtual server.
 - Define policies, service classes, and classification rules for the workload, and schedule activation of the appropriate policy. Note that z/OS WLM restricts the length of service class names to 32 characters, so you need to observe this restriction when you define the names of zManager service classes that you want to map to WLM service classes.
 - Configure and activate a guest platform management provider on virtual servers that support the tiers of a multitier application. Also enable application response measurement (ARM) for the IBM middleware running on those virtual servers.
- Note: Depending on the configuration requirements of the multitier applications that your company uses, you might need to configure guest platform management providers and enable ARM for middleware not only on the virtual servers running z/OS, but also on z/VM guests and virtual servers on a POWER blade.

For z/OS images running on virtual servers

- For each zManager workload that contains virtual servers running z/OS:
 - Decide whether corresponding WLM service classes are required for the service classes within the URM workload policy. If so, define the new WLM service classes that set only response time goals for URM work requests. These response time goals apply to all DB2 distributed data facility (DDF) requests that originate from ensemble virtual servers that are classified within the URM workload performance policy.
 - Optionally, define one WLM service class as the default service class for work associated with a URM workload policy.
- To map URM service classes to their corresponding WLM service classes, modify the classification rules for the EWLM subsystem type. Use only the ESC work qualifier in the classification rules.
- Optionally, specify WLM report classes to monitor the performance of URM work requests separately from local z/OS work.

Unified Resource Manager is a.k.a. zManager