



# IBM Maximo Asset Management Report Performance Analysis Guide

The collage displays several key reports from the IBM Maximo Asset Management system:

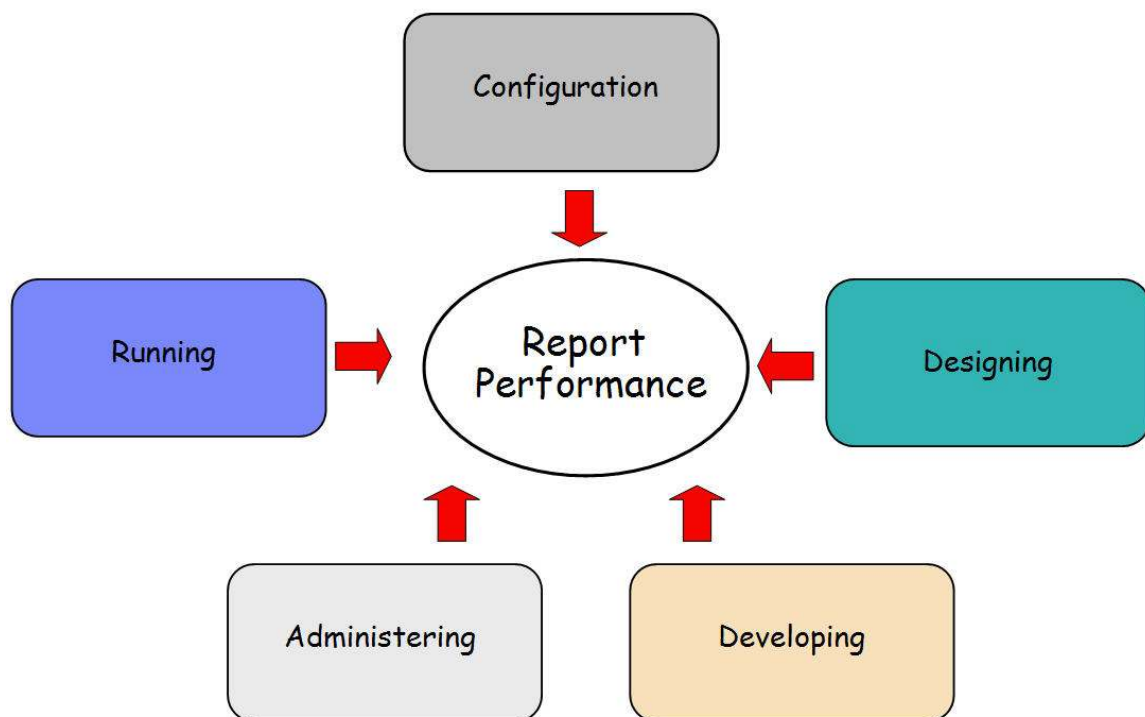
- SLA Details:** A table showing SLA information for a specific asset, including status (ACTIVE), incident type, customer, and classification.
- Login History:** A line graph showing the number of concurrent logins over time, with a table below listing user details.
- Asset Availability:** A bar chart and table showing the availability percentage of assets over time.
- Software Usage Summary:** A pie chart titled 'Software Usage - By Vendor' showing the distribution of software usage across different vendors.
- Service Level Exception:** A report detailing service level exceptions, including the asset name, status, and date.
- Work Order Details:** A detailed view of a work order, including task descriptions, dates, and status.
- OEE% by Equipment:** A bar chart showing Overall Equipment Effectiveness (OEE) percentages for different pieces of equipment.
- Item Availability:** A table showing the availability status of various items, including item ID, description, and location.

V7 Report Performance Overview .....	3
Configuration .....	6
1. Configuration: Clustering.....	7
2. Configuration: BIRT Report Only Server .....	10
Configuration Recommendations and Notes:.....	17
3. Configuration: Replicated, Reporting Databases .....	18
4. Configuration: Hardware.....	19
Designing.....	22
1. Design: Analysis Options .....	23
2. Design: Format Impacts.....	25
Developing.....	29
Administration.....	30
1. Administration: Record Limits .....	31
2. Administration: Queuing .....	33
3. Administration: Performance Tab .....	36
4. Administration: View Report Processing.....	37
Execution .....	38
1. Execution - Scheduling.....	39
2. Execution: Schedule Only Reports .....	41
3. Execution - Schedule Only Reports - Reserved Processing Time .....	45
Report Schedule Limits.....	49
Set Report Schedule Limits.....	50
Managing Report Schedule Limits.....	52
4. Execution - Report Usage.....	53
Additional References.....	55
Revision History .....	56
Notices.....	57
Trademarks .....	58

## V7 Report Performance Overview

To respond to today's dynamic Business Environment critical business information needs to be immediately available. This business information can come in a variety of formats, and is often required as a formatted business report, or Enterprise report.

In order to meet these increasing information needs of users, this document will review Report Performance and the components that impact it. These five components - Configuration, Designing, Developing, Administering and Running - must each be synchronized in order to maximize report performance. If any of the five components is not analyzed and optimized, the entire reporting process will not be optimized.



The five specific reporting components below are discussed in this document. Evaluating each of these components will enable you to optimize report performance in your unique environment.

\*Note: This document applies to the Maximo® Base Services 7.x Releases.

### Configuration

Configuration involves not only the architecture of the entire System, but also its hardware specifications (CPU, Disk Space), Network Connections and Directory Service (LDAP). This document will focus on the areas directly impacted by performance including: Clustering, Report Only Server, Replicated Databases and Reporting Hardware Requirements.

### Designing

The design of a report greatly impacts a reports performance. A simple list report with a single grouping could take just a few seconds to execute, whereas a complex analysis report containing over ten subreports, multiple graphs and groupings could take much, much longer.

Additionally, whether or not a report is the best mechanism for the data analysis is another area that needs to be discussed. Often times, users only need to know the status of a particular business event, like 'Number of Work Orders Outstanding', 'Invoices Past Due', or 'SLA Compliance'. These are visual metrics whose values may be better created and maintained via the KPI Manager.

### Developing

The development of a report is very critical due to its sql statement. An optimized sql statement enables a report to execute at its peak. However, using non-optimized sql statements with incorrect table joins can cause a report to not only return incorrect data, but execute inefficiently consuming critical system resources.

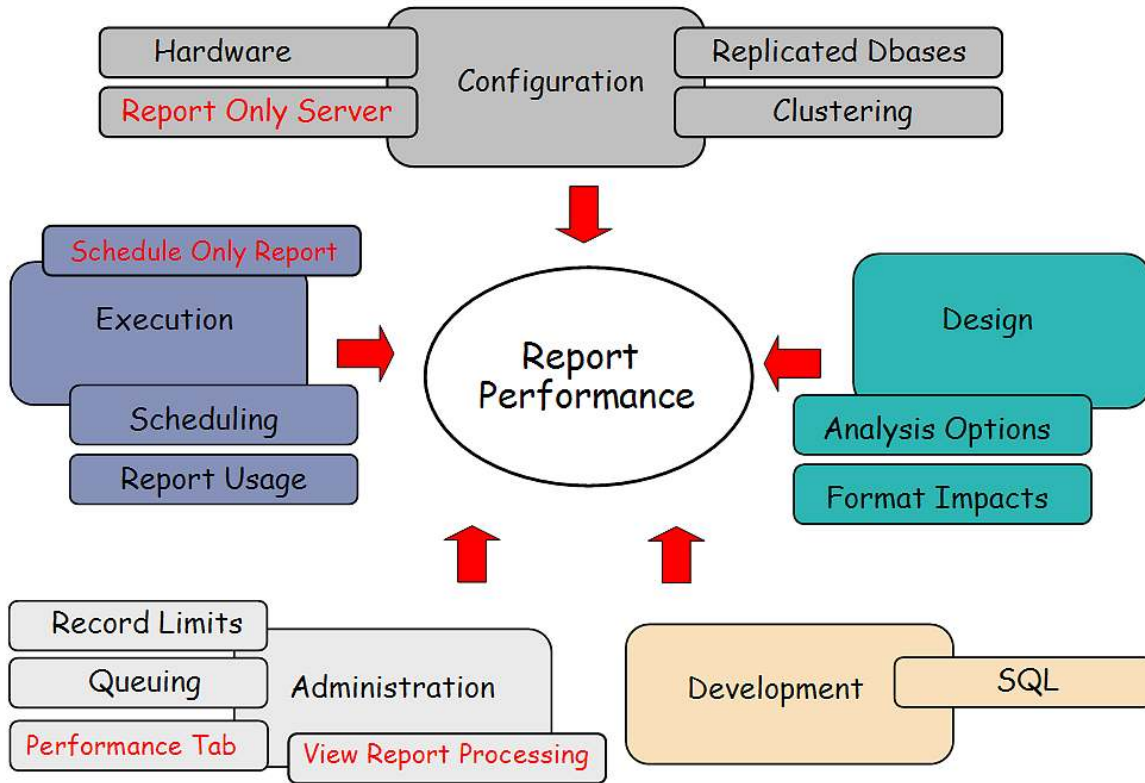
### Administering

Administration of reports is very critical to their performance success. Specific items that administrators can focus on are Enabling Record Limits to prevent end users from inadvertently executing reports against large data record sets, and enabling report queuing. Report queuing controls the number of reports that can be executed at any time to insure optimal performance across the V7 Server. Additionally, administrators can use tools - like the Performance Tab and Viewing Report Processing - to monitor report performance.

### Running

End Users can optimize report performance by scheduling complex, batch reports to execute at off-peak hours. Additional features are available to users to manage their own reports to insure optimal performance.

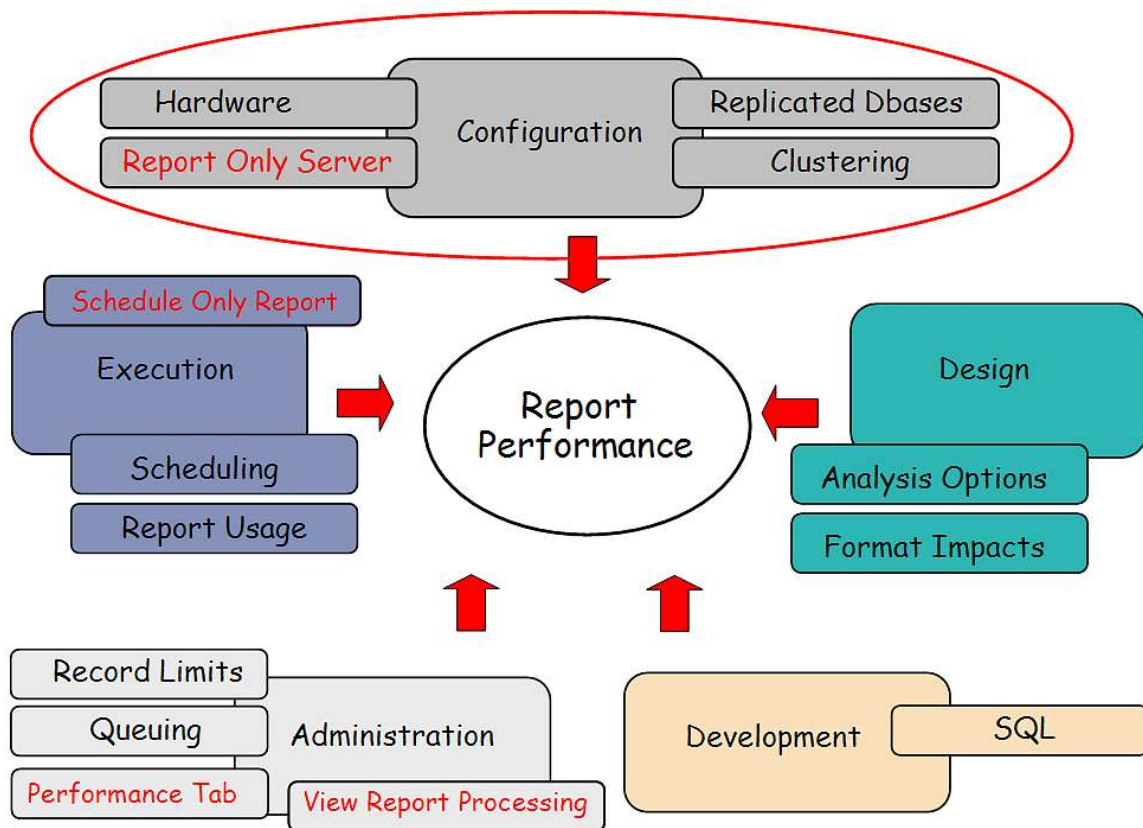
For each of the five components, segment(s) of each will be explained in more detail below.



Note: The items with text in red: Report Only Server, Schedule Only Report, View Report Processing and Performance Tab - were introduced in Version 7.1.1.4 in December 2008.

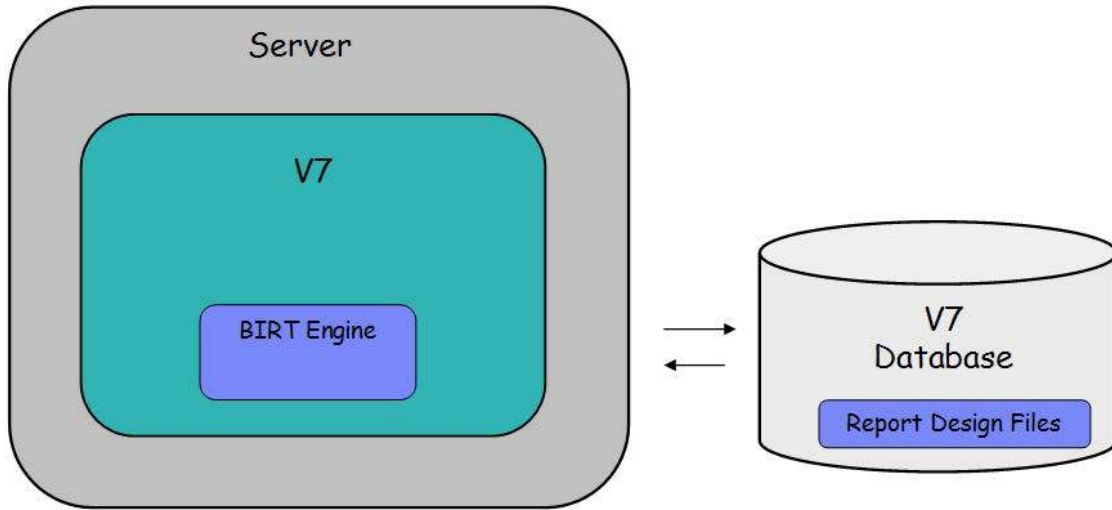
## Configuration

This section will focus on four key areas of Configuration for Reports: (1) Clustering (2) Report Only Server (3) Replicated Databases and (4) Hardware.

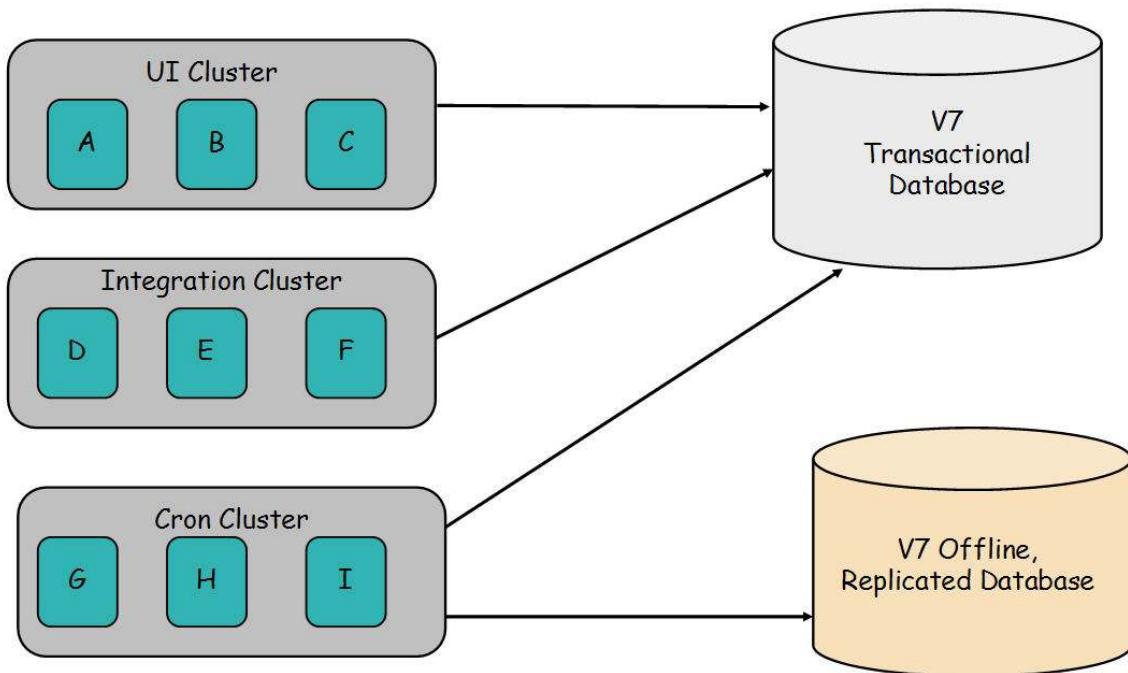


## 1. Configuration: Clustering

In a very simple development or test environment, the V7 application may be deployed on a single application server. This environment is then configured to a single V7 database.



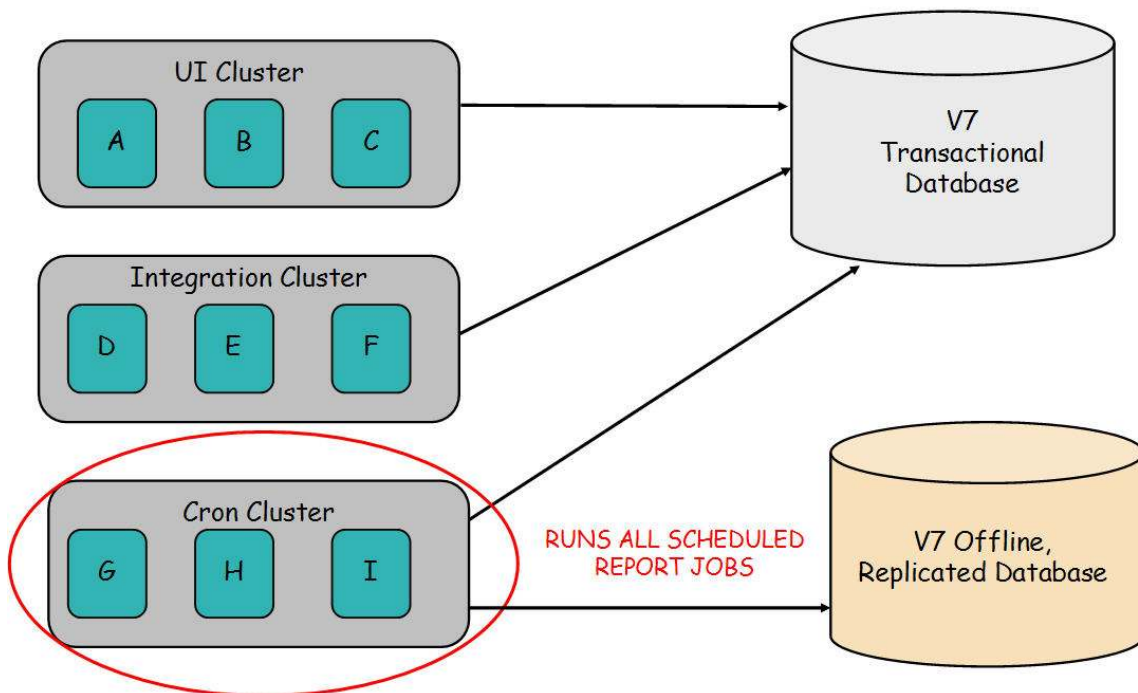
Advanced Server Configurations are required in production environments to enable the system to scale effectively and enable the best performance for all users. At a top level, three types of clusters are recommended which are UI, Integration and Cron Clusters.



UI Cluster: The UI cluster is intended for users to access the applications from a Web Browser. This cluster is typically set up for load balancing so that users accessing this browser only use one URL to access V7 that may be running on any one of the servers within this cluster.

Integration Cluster: This cluster processes integration messages from message queues, and moves messages into the queues by using HTTP Post, Web Services and EJB Technology.

Cron Cluster: The Cron cluster processes scheduled jobs. These scheduled jobs can be related to integration and reports. This section will focus on the Cron cluster, and enabling it to executing all the scheduled report jobs.



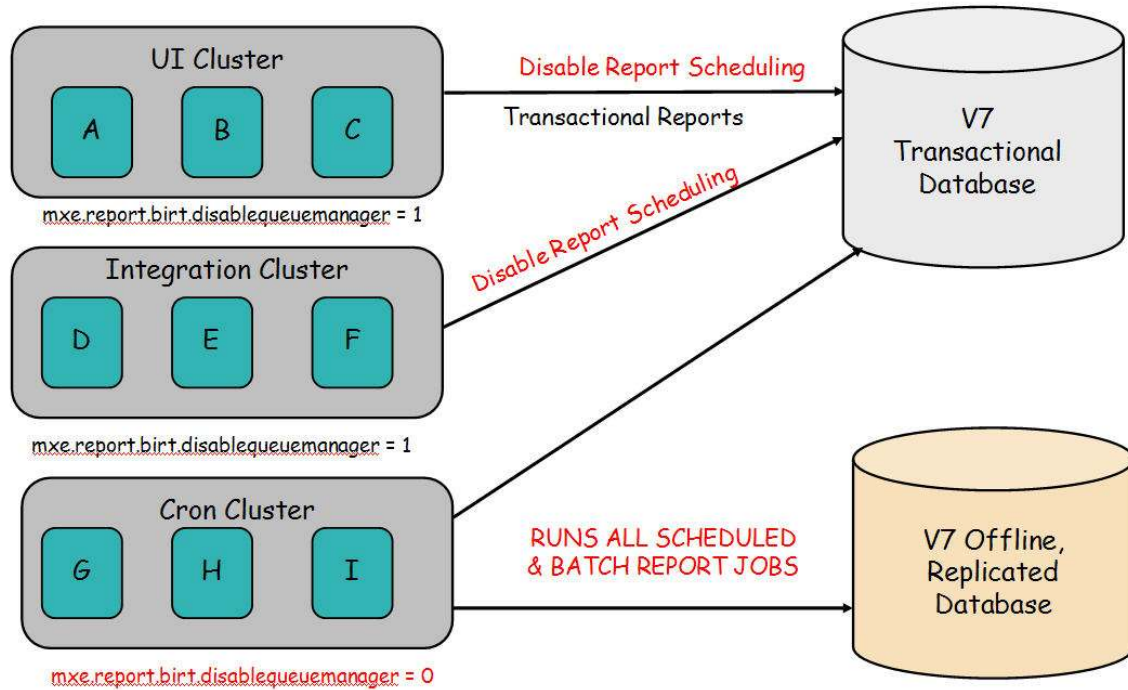
To optimize performance in a clustered environment, it is recommended that all scheduled report jobs be executed from the Cron cluster. This will divert the scheduled report load from the UI cluster to the Cron cluster. This diversion will increase the response time users will experience in the UI cluster.

Users in the UI cluster can still execute immediate reports to receive the immediate report information they need. However, the more complex, analysis type reports that are scheduled will be off-loaded to the Cron cluster of servers.



To enable report scheduling to only occur within the Cron Cluster, a property value needs to be configured. The property value is `mxe.report.birt.disablequeueemanager.*`. This value should be set to 1 for any server where report scheduling is not enabled.

If the queue manager is disabled (value set to 1) scheduled reports will not be executed on the server.




More details on clustering and enabling its setup are contained in the V7 System Administration Guide, in the System Configuration Chapter.


## 2. Configuration: BIRT Report Only Server

Beginning with the 7.1.1.4 release, you have the option to configure their V7 Environments to have a BIRT Report Only Server (BROS). Enabling BROS will balance report load processing and can improve overall system performance. The BROS will be utilized for report processing - regardless of what clustered server the user may be on.

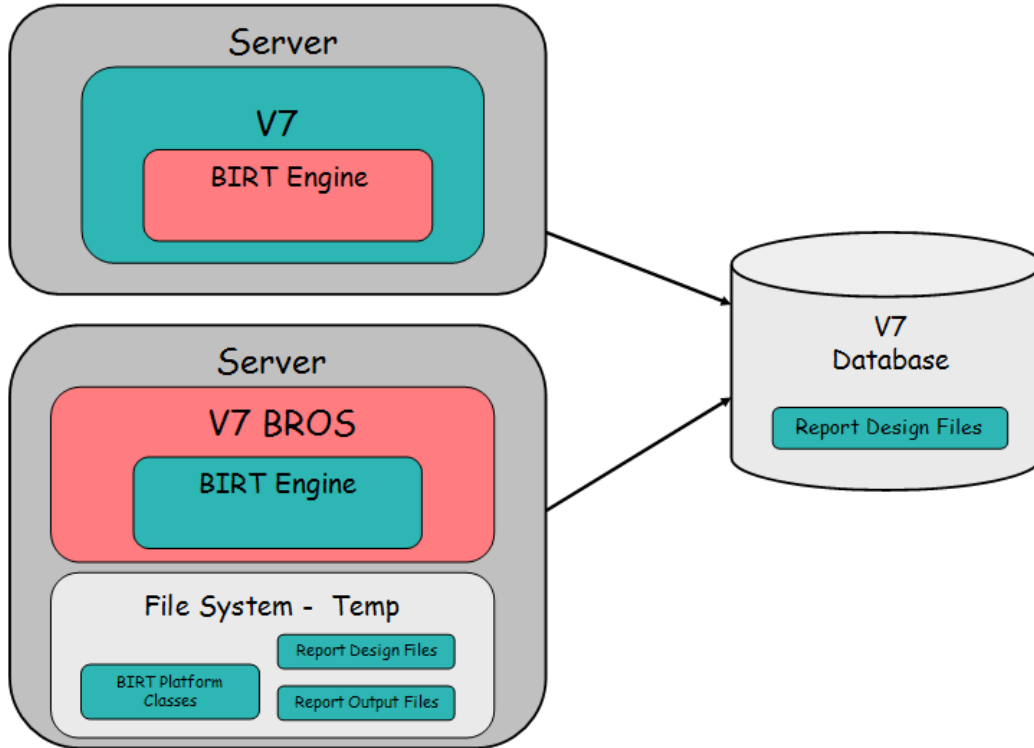
To highlight this functionality, imagine a very simple V7 environment. In this environment, the V7 Server is deployed within an application server. This deployment automatically installs BIRT through the ear file process. When a user performs any V7 action or executes a report, both processes occur within this same server.



You can extend this environment to configure an additional server to handle reporting processes. This will remove the performance impact by off loading reporting processes to a different server.



With the BIRT Report Engine on a separate server, the dynamic hardware requirements required for the report temporary files will only be required on the BROS. These temporary file requirements include the copies of the report design files, the report output files and the BIRT platform classes.



To enable the BROS Configuration, a property file will be used to identify the URL of the report server. Its value is: `mxe.report.birt.viewerurl*`

This property setting is not required.

- If its value is not defined, reports will execute from the V7 Maximo Server.
- If the property is defined, reports\* will only execute from the Report Only Server, and/or the Cron Task Server (If Report Scheduling is off-loaded to a Cron Cluster).

\*Note: Certain report functionality will continue to execute from the UI server, even when a BROS server is configured. This functionality includes:

- A. Direct Print Reports
- B. Direct Print with Attachments Reports
- C. Query Based (Ad Hoc) Reports as they are being created. Once a QBR report is saved, it will execute from a BROS server if configured.

The diagram below shows an example of how this could be configured. In this example, there are two servers:

- V7: UI server used for execution of V7 Applications
- V7 BROS: Used only for report processing of V7 reports

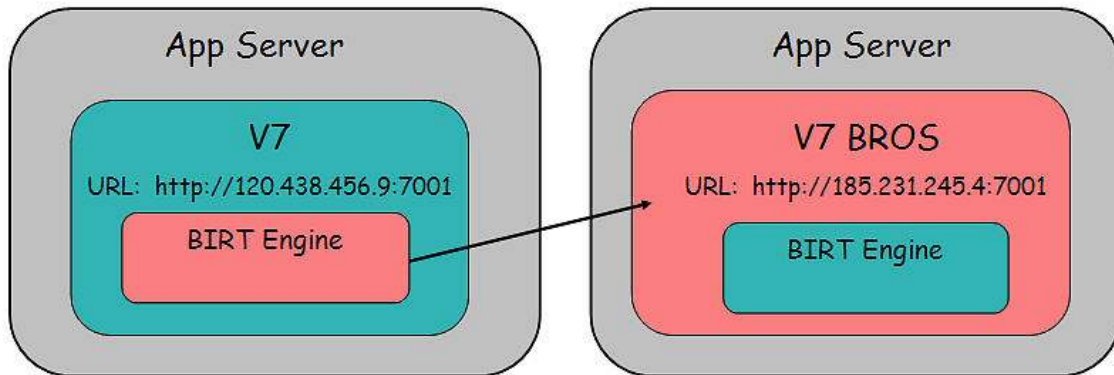
To enable this configuration, two property settings are required.

The first global property setting points to the location of the BIRT Report Only Server, `mxe.report.birt.viewerurl`.

This value is set to the location of V7 BROS: `http://185.231.245.4:7001`

The second property setting determines if report scheduling should be executed from the server. This is property setting `mxe.report.birt.disablequeue manager`. In this example For V7, this value is set to 1 (Yes) because no report scheduling should occur on this server.

For V7 BROS, this value is set to 0 (No) because report scheduling should occur on this server.



`mxe.report.birt.viewerurl = http://185.231.245.4:7001`  
(Enables reports to execute from BROS Server)

`mxe.report.birt.disablequeue manager = 1`  
(Disables report scheduling from V7 Server)

`mxe.report.birt.disablequeue manager = 0`  
(Enables report scheduling from BROS Server)

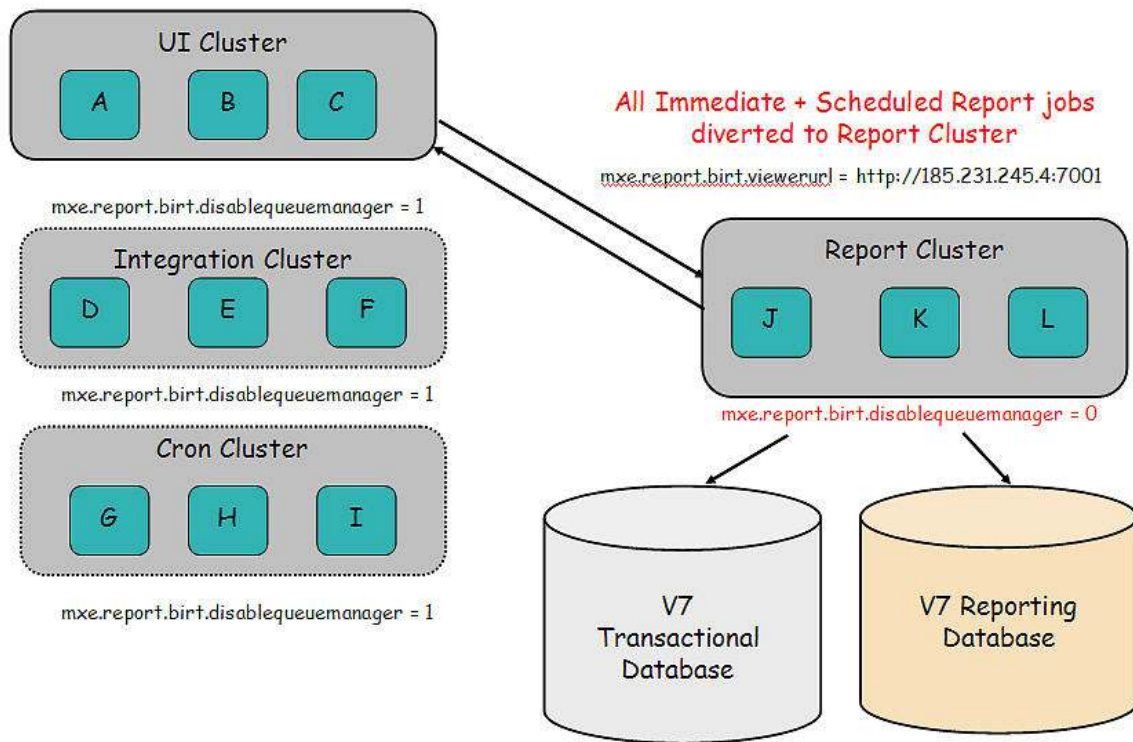
#### Notes:

1. The hostname in the Version 7 URL cannot be the same as the hostname in the BROS URL. For example, if the user logs into Maximo as <http://abc.com:9080/maximo>, the hostname is "abc.com". The BROS value, `mxe.report.birt.viewerurl`, *can not* also use abc.com.
  - If the same hostnames are used, the Maximo UI Session will become invalid, and the user will be logged out.
2. When a BROS is configured in a your environment, the BIRT Engine will still be deployed on the V7 Maximo Server. However, it will not be used.
3. The examples used here are for non-SSL Enabled Environments. If your environment is SSL enabled, be sure to use https in your url values.

With the BROS configuration, you can enable a number of different clustered server arrangements. One arrangement that would be recommended is to extend the three cluster configuration (UI, Integration and Cron) to a four cluster configuration, which would include:

1. UI
2. Integration
3. Cron
4. Report

This configuration could improve overall system functionality by removing the report processing load from the UI Cluster. All report jobs - both immediate and scheduled - would be diverted to the Report cluster. This configuration is shown in the diagram below. Each of the four clusters contains three servers noted by an alphabetical character.

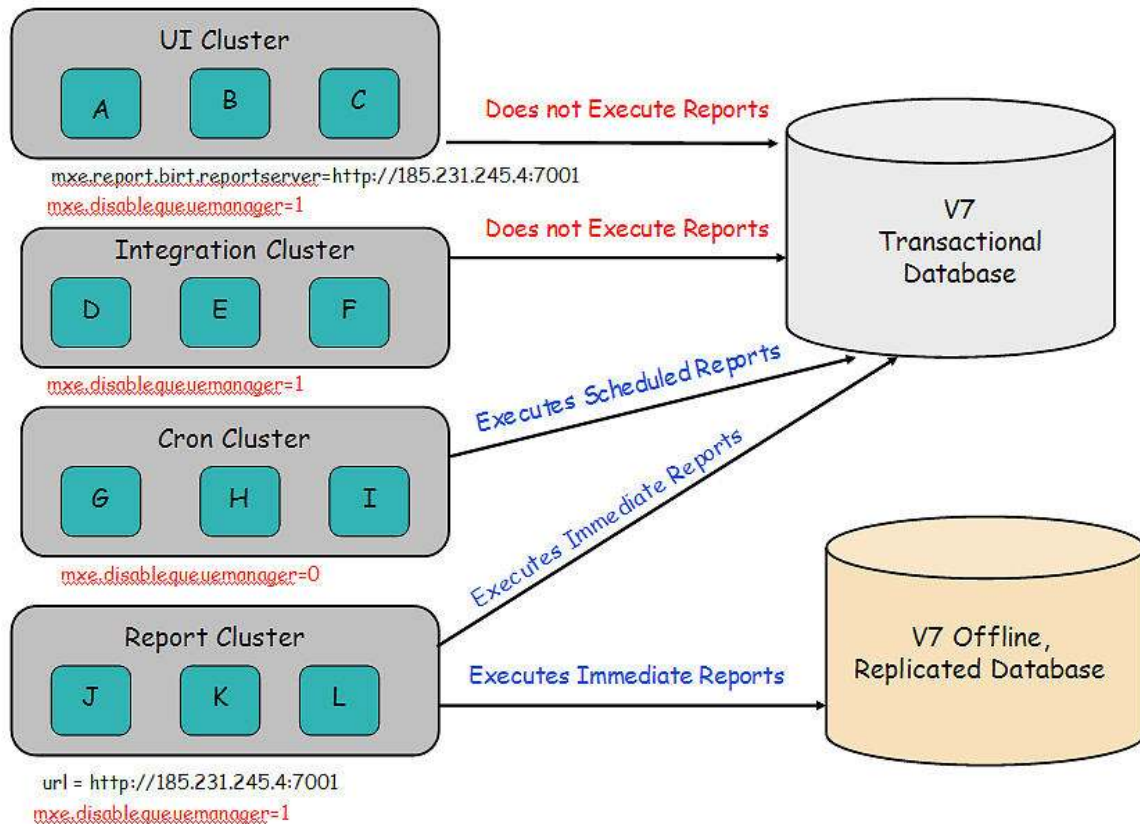


In order to enable report jobs - both scheduled and immediate - to run off the Report Cluster, the Report Queue Manager has to be disabled from the other servers. Disabling the Report Queue Manager prevents scheduled reports from executing on that server. (More information on the Report Queue Manager is contained in the V7 Report Feature Guide.)

To disable the Report Queue Manager, the property `mxe.report.birt.disablequeue manager` must be set to 1. In this example, the UI, Integration and Cron Servers all have this property set to 1. Additionally, in this configuration, the server where maximo is deployed (UI Cluster) would have a defined URL for the `mxe.report.birt.viewerurl`. The URL would point to the location of the BROS.

Multiple servers can be used for report servers in a clustered environment. In the diagram above, three Report Servers (J, K and L) are included in the Report Cluster. However, to enable this, only one reportserver URL is required. This report server URL essentially points to the load balancer, which directs the report job request to any number of report servers defined within the cluster.

In another example of a clustered environment, four clusters are still used. However, in this case, the report scheduling will be done in the Cron Cluster, and the immediate report processing will be done in the Report Cluster.



This will again remove the load from the UI cluster, and place the report processing in both the Cron and Report cluster. This arrangement would be ideal if you process a large number of immediate report jobs as the report cluster would fully focus on processing immediate jobs only.

In order to enable the scheduled report jobs to execute from the Cron Cluster only, the Report Queue Manager has to be disabled from the other servers. Disabling the Report Queue Manager will prevent scheduled reports from executing on that server.

To disable the Report Queue Manager and hence, report scheduling, the property `mxe.report.birt.disablequeueuemanager` must be set to 1 in the UI, Integration and Report Servers. Only the Cron Servers will have the `mxe.report.birt.disablequeueuemanager` set to 0 to enable report scheduling on its server. \*



**\*Configuration Set up Note:**

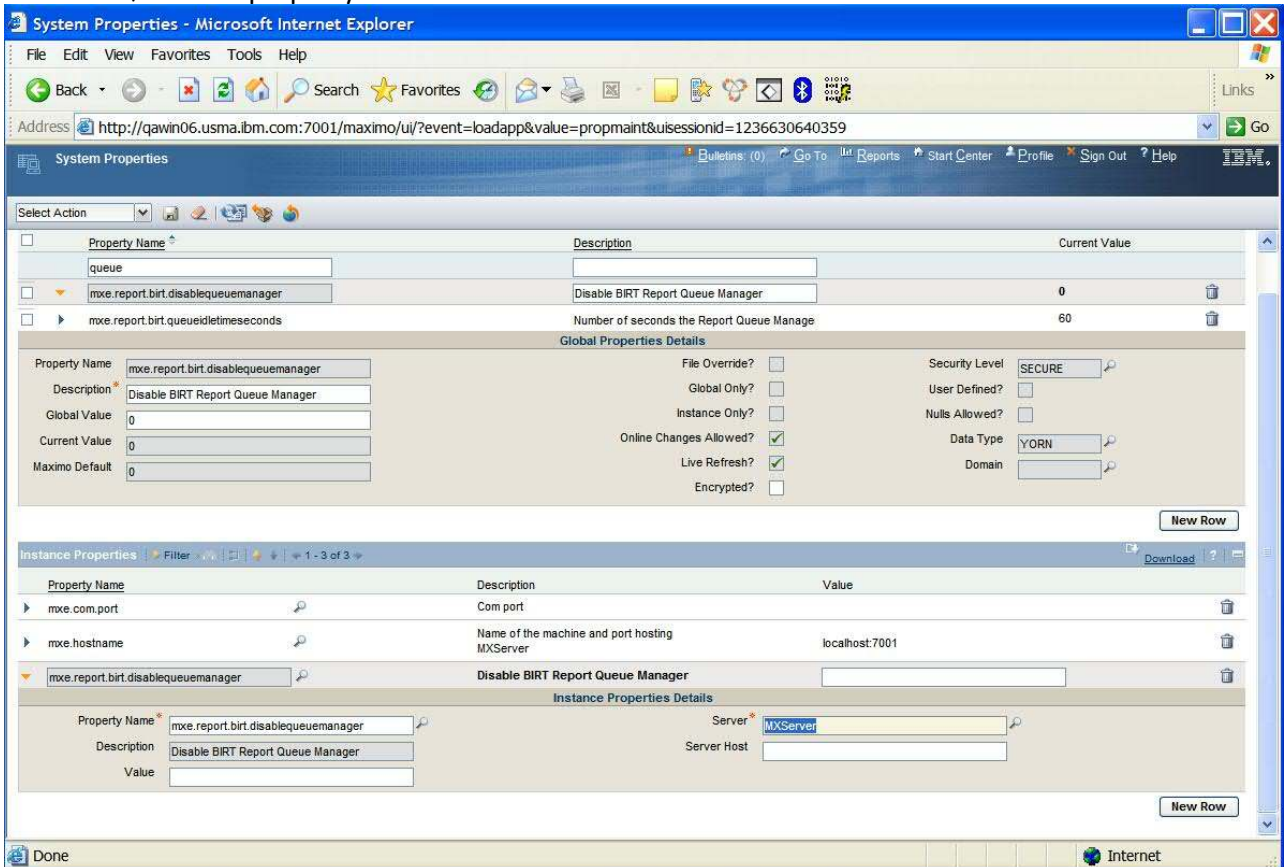
The queue manager property file can be overridden at the server instance level if needed. Information on how to do this is contained in the *Maximo System Administration Guide* on page 234. *mam71\_sys\_admin\_guide.pdf*. Its text is also copied here below, along with a screen shot on how to enable this.

## Instance Properties

An instance property is defined and associated with a specific system server. For example, you can configure the system property, `mxe.crontask.donotrun`, to be an instance-specific property. To perform this configuration, access the System Properties application. In the Instance Properties table window, associate the `mxe.crontask.donotrun` system property with a specific server (for example, `MXServer1`) and a value applicable only to that server (for example, a value of `BBCron`). As a result, the Bulletin Board cron task (`BBCron`) does not run on `MXServer1`; however, `BBCron` can run in another system instance, such as `MXServer2`.

When you create both a global value and an instance-specific value for the same property, the instance-specific value takes precedence.

To do this, add the property name and server as shown in the screen shot below.



## BROS Authentication

When a BROS is configured, authentication information needs to be passed from the V7 Server to avoid re-login into the BROS that receives the request to process the report.

This authentication is not required when the report server is on the same server as V7 as the two share a session.

To enable authentication for the separate Report Server, tokens will be used. This will enable authentication whether the configuration uses the application server authentication or V7 authentication. The top level steps describing this authentication are:

1. When a report is executed, a random token is generated.
2. This token is stored temporarily in the database table, MAXTOKEN, and is passed to the BROS.
3. The report process servlet(s) receives the token and verifies it with the database.

If the verification is successful, an authenticated session is populated in the HTTP Session, and the token is removed from the database.

Subsequent report requests will then check for the authenticated session information and process the request.

Use cases may result where a token is generated and stored in the database, but the report process servlet does not get a chance to process the report request. In this case, token values can accumulate in the database. To prevent this from happening, token records will be cleaned up via a configured property setting, called `mxe.sessiontoken.timeoutseconds`. By default, this value is set to 180 seconds, or 3 minutes.



## Configuration Recommendations and Notes:

1. The BROS server does not support application server security. As noted above, BROS utilizes token based functionality for authentication.

However, LDAP is supported with BROS and Maximo if the domain name of BROS and Maximo is the same. For example if maximo's url is abc.xyz.com and the BROS url is def.xyz.com - then, this configuration will work as they have the same domain name.

2. Enabling the BROS Configuration does not change the install or the V7 Ear file deployment process. The same configurations are deployed on all servers. For example, the entire V7 application is still deployed on the BROS Server - however, only the reporting functionality is be enabled.

3. When the report server is deployed inside a V7 Server, it is recommended that no more 5 reports run concurrently at any given time. This value is defined via the property file, mxr.report.birt.maxconcurrentrun.

The value of 5 is based on the V7 Server which is scaled for 50 concurrent users per server. Estimates indicate that that 10% - or 5 users - will be running reports concurrently. Therefore, the value of 5 is used.

However, when a BROS is configured, the prior maximum number of 5 concurrently running reports may no longer apply. In this case, it is recommended to increase the property value, mxr.report.birt.maxconcurrentrun, to '2\* Number of CPUs' of the Server.

4. Also, multiple clients have asked that since V7 is scaled for 50 concurrent users per server, does this mean that the BROS can also handle 50 concurrent report users?


For example, a client has 500 V7 Users. He uses 10 V7 Servers to handle the 500 users. Assuming the standard 10% of those users at any time are concurrent report users, can 1 Report Server handle the 10% or 50 Concurrent Users running reports from these 10 different V7 servers?

The answer to this is not straightforward as it depends on the types of reports that are being run. Reports have a huge range in complexity, from a very simple list report to a very complex Maintenance Cost Rollup report that encompasses numerous subreports and hierarchy levels. Therefore, this answer really depends on the average complexity of reports that are being executed.

### 3. Configuration: Replicated, Reporting Databases

You may want your reports to execute from a separate, reporting database. This separate, external database may be a snap-shot of their production, transactional database, and is sometimes referred to as an external, replicated or reporting database.

Using reporting databases can reduce the load from your live, transactional database and lead to performance improvements. Additionally, if these reports are scheduled, additional processing time can be off-loaded from the V7 Server. These performance improvements can be especially noticeable for complex reports.



This type of configuration can be enabled two different ways:

1. All reports can be configured to execute against a reporting database,
2. Or a portion of the reports can execute against the reporting database.

To enable these two different types of configuration, the Reporting Database (Datasource) needs to be registered in the Report Administration application. Then, if only a subset of reports will execute from the Reporting Database, additional steps will be required to modify their individual report design files within the BIRT Report Designer.

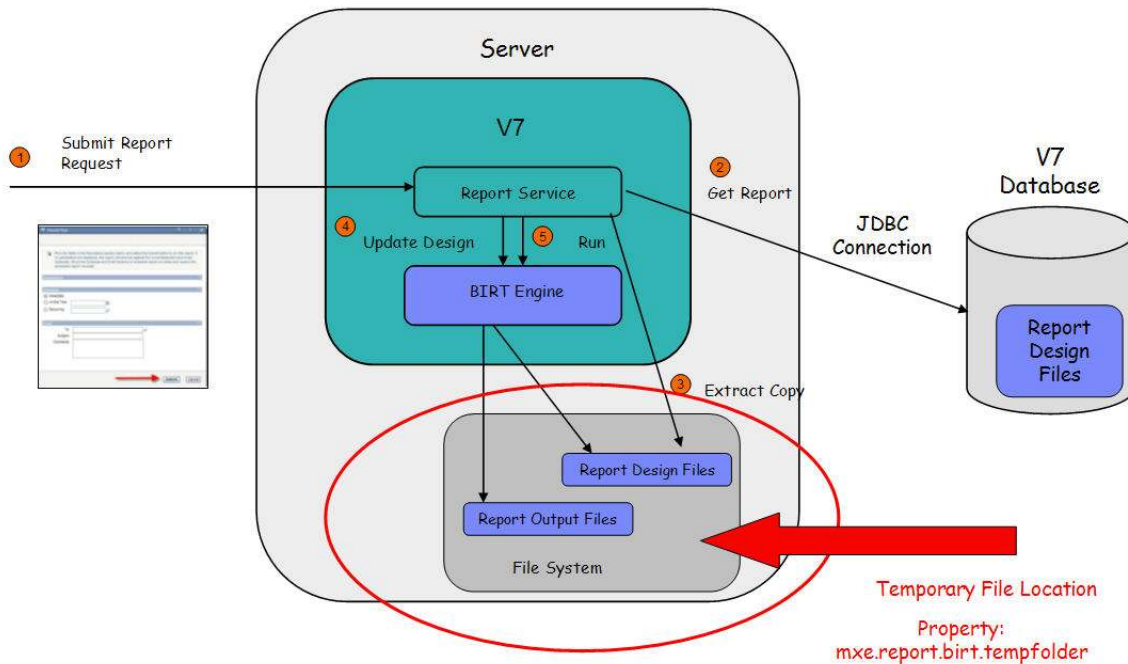
The steps on how to configure each scenario are contained in the 'Enabling Secondary Database Configuration for Reporting' Document which is posted on IBM's Support Website. Its Document Reference Number is 1304936.

#### 4. Configuration: Hardware

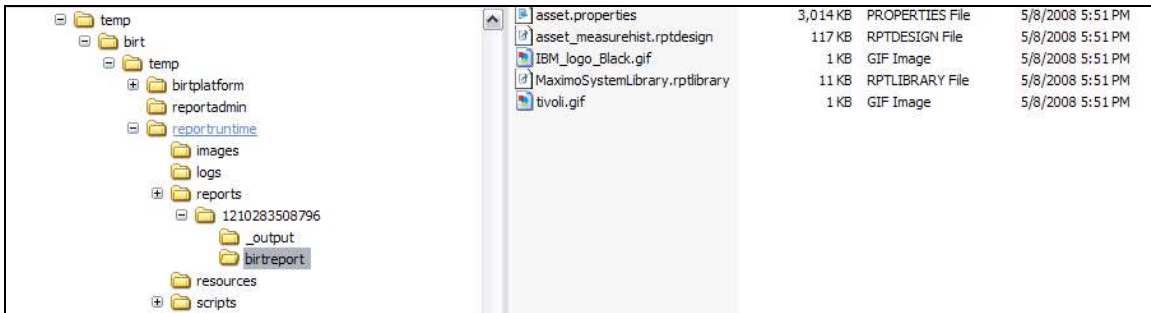
As noted in the V7 Reporting Feature Guide, temporary disk space is required on each V7 Server where you will be executing BIRT reports. This disk space is needed to hold temporary copies of the BIRT design files and graphic images, along with the BIRT Platform Classes.

The BIRT Platform Classes are required in the temporary location for initializing the BIRT Report Engine. These files are approximately 24 MB, and will always be required.

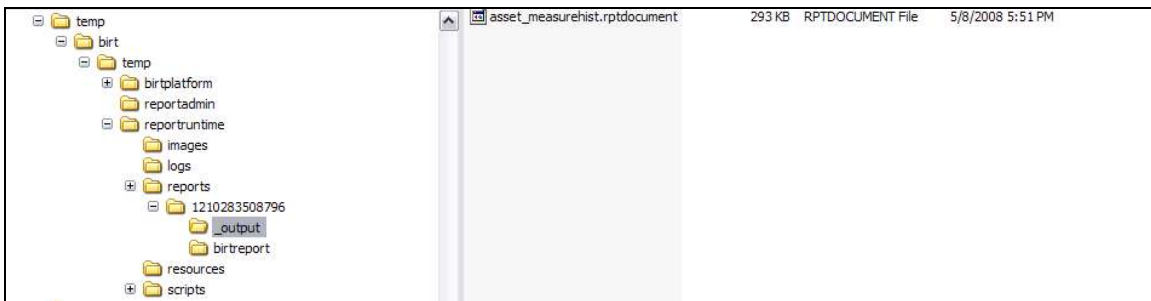
The report design files and the report output files are also stored in the file system's temporary directory as highlighted by the red arrow below. These temporary files are required when users are actively running reports.



The report design files are the temporary copies of the user's requested report from the database. This will include not only the .rptdesign file, but also any libraries or resource files it uses (ex. Property files and Resource .gif/.jpg images.) An example of a report stored in the birtreport subfolder is shown here:



The report output files are the actual reports that are generated based on the user's individual request. These files have the extension of .rptdocument, and are stored in <c>temp\birt\temp\reportruntime\reports\\_output



The amount of file space required on your V7 Environment for the temporary report design and output files depends on your individual business environment. The file space required depends greatly on the:

1. File size of the report design and its output
2. Whether or not the report contains graphic images or not
3. How many reports are being concurrently run by users
4. How long the V7 session is kept open

These temporary files are retained until a cleanup mechanism deletes them. This mechanism is invoked according to certain business rules, including:

- A. If the user runs the same report multiple times, the previously run report information from the temporary folder will be removed.
- B. If the user signs out of V7, then all the temporary report information related to all the reports that the user has run from the browser will be removed.
- C. If the user session times out (which is by default, 30 minutes), then all the temporary report information related to all the reports that the user has run from the browser will be removed.
- D. If the server is restarted, then all the temporary report information related to all the reports that any user has run from the browser will be removed.

Next, an example will be given which will highlight the variability of this number. In this example, there are 500 V7 Users. Of those 500 Users, 10%, or 50 Users are concurrently running reports.

Each of the 50 users executes a report that has an output of 1MB.

*For these 50 users, 50MB of disk space is needed. (50 Users \* 1MB)*

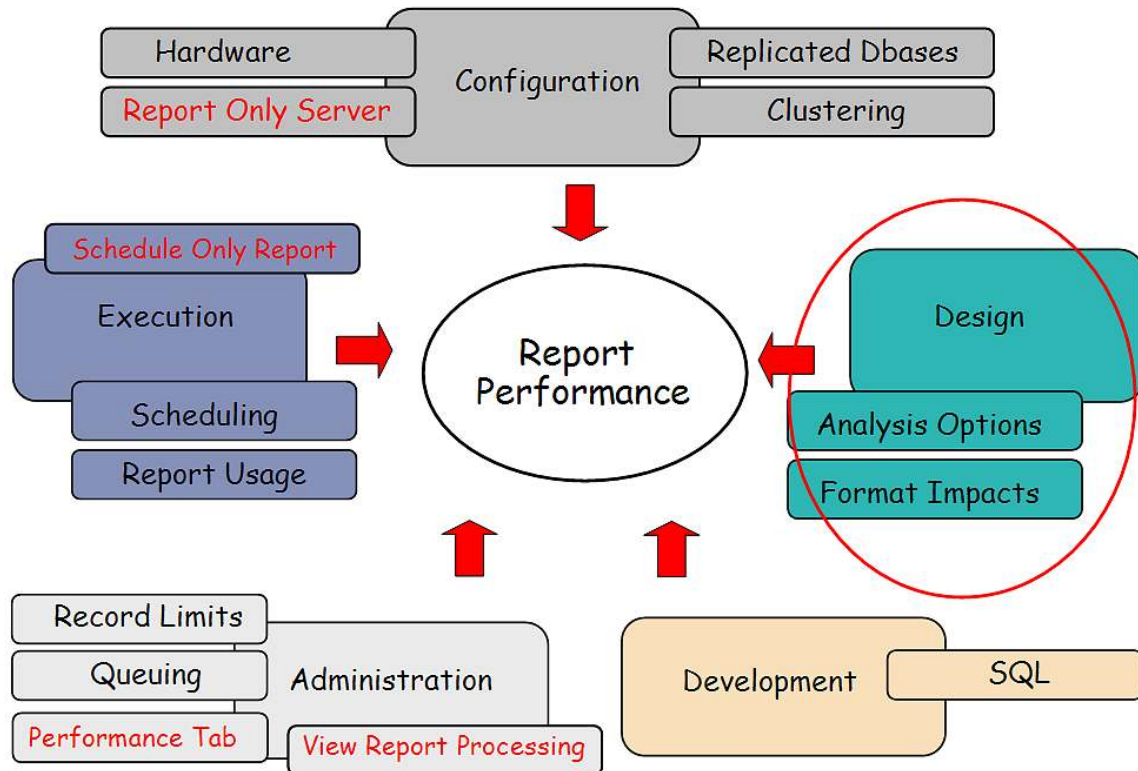
Now, imagine each of the 50 users each run 5 reports.

*For these 50 users, 250MB of disk space is needed. (50 Users \* 5 Reports \* 1MB)*

As this example shows, the amount of temporary file space you required will vary depending on the number of active report users you have along with the report file sizes.

## Designing

This section will focus on some of the areas within the Report Designing Process which impact reports including (1) Analysis Options and (2) Format Impacts.



## 1. Design: Analysis Options

Within the various applications, there are a variety of mechanisms available to analyze the powerful data that is generated. These mechanisms present the data in a variety of formats, so you can quickly get the information needed to make the best business decisions possible.

Depending on the individual business case and user, you may want to evaluate using different Data Analysis Options than reporting in certain cases. For example, one of your users may need to know the Number of Work Orders Overdue each day when he logs in. This information is critical for him to meet his job objectives, and he wants to quickly see this data. In this case, a KPI visually displaying the Number of Work Orders Overdue on the user's Start Center may be a better fit to convey this information.

When a report is created, it must first be designed, developed, tested and maintained. Depending on the complexity of the report, this could take a few hours, a few days or a few weeks. Then, once those reports are available, they will require system resources to execute. Again, depending on the complexity of the report, those system resources could be minimal or significant.

The latest data analysis features available in Version 7 are listed below. These features enable different forms of data analysis, which can quickly and dynamically display data in any applications. Additionally, they can be considerably less time consuming to enable than creating individual custom reports, and therefore should be carefully considered.

The different functionality formats include:

*QBE - Query By Example.* Using your application's filter and/or query, you can immediately download your results for additional analysis in Microsoft Excel.

*RS - Result Sets.* Using an application's query, you can enable a set of fields or graph for display on the Start Center.

*Application Export.* Capitalizing on the power of an object structure, you can export multiple database objects and attribute fields using your application's filter and/or query.

*KPI - Key Performance Indicators.* Visual indicators displaying status against predefined targets.

*QBR - Query Based Reporting.* Version 7's version of Ad hoc reporting where your users create their own reports on the fly from within the various applications.

*OR - Operational Report.* Often referred to as transactional reporting, these are the day to day detail reports your users require to complete their business tasks.

*SR - Strategic Report.* Enable viewing of data in varying perspectives thru the use of complex graphs, in depth calculations or scenarios.



For additional details on these data analysis options, reference the V7.1 or V7.5 Upgrade Planning Guide referenced at the end of this document.




## 2. Design: Format Impacts

The specific format designed into a report has a significant impact to its performance. Reports that contain multiple graphs, groupings and controls will require a longer time to execute than simple listing reports.

To highlight the impact of formatting, a variety of tests were performed using the V7 Out of the Box Reports that are delivered. This testing was done to quantify the performance impact of graphs, groupings and controls that you could potentially see.

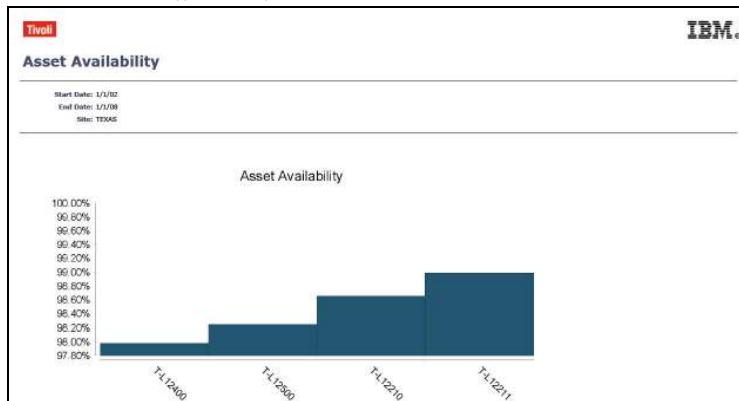
The chart below shows the largest performance impact seen by reports that included graphics, followed by those with a large number of controls and multiple groupings. While over 300 reports were executed to determine these results, your individual results may vary depending on the specific designs of the reports tested.

Each of these tests was performed in the same V7 Environment. A variety of V7 Users executed these reports, which were accessed from a variety of applications, and from a variety of actions within the application, including application toolbar access and the action menu.



## Graphs

Graphs are extremely useful in quickly conveying data results to users. Their visual representations often immediately highlight results, so users do not have to manually analyze the data themselves.




To quantify the impact of graphic rendering on report performance, five different OOB Reports were analyzed, including Asset Availability, Inventory ABC, Asset Measurement History, Software Usage and Service Target Compliance. These reports included a variety of graph types, including line, bar and pie charts. Each report included 1 graph, except for the Service Target Compliance Report which contained 3 Graphs.

To perform the test, each report with the graphic was executed 10 times. Then, the graph from each report was removed, and the reports were then executed 10 more times. These results were captured in the REPORTUSAGLOG Table.

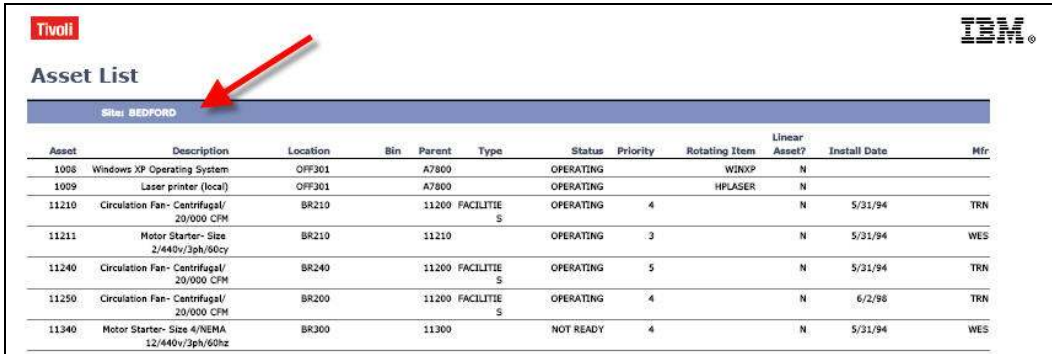
The results below show that on average, reports with no graphs ran 38.8% faster than those same reports that contained graphs.

These results highlight the significant impact of graphs on performance. This data can be used to validate that you may only want to use graphs in reports that are analysis type reports. Analysis reports are typically scheduled or executed against non-Transactional Databases. Including graphics in transactional reports that need to be executed immediately and consistently throughout the day will have a negative performance impact.



## Groupings

Groupings are used in reports to separate data. For example, if you are analyzing an asset list report, the assets may be grouped by Site and/or Location. This enables the user to break down the information. So - instead of evaluating 100 Assets, the user can break it up into 40 Assets in Site Bedford, 25 in Site Texas and 35 in Site Buffalo.



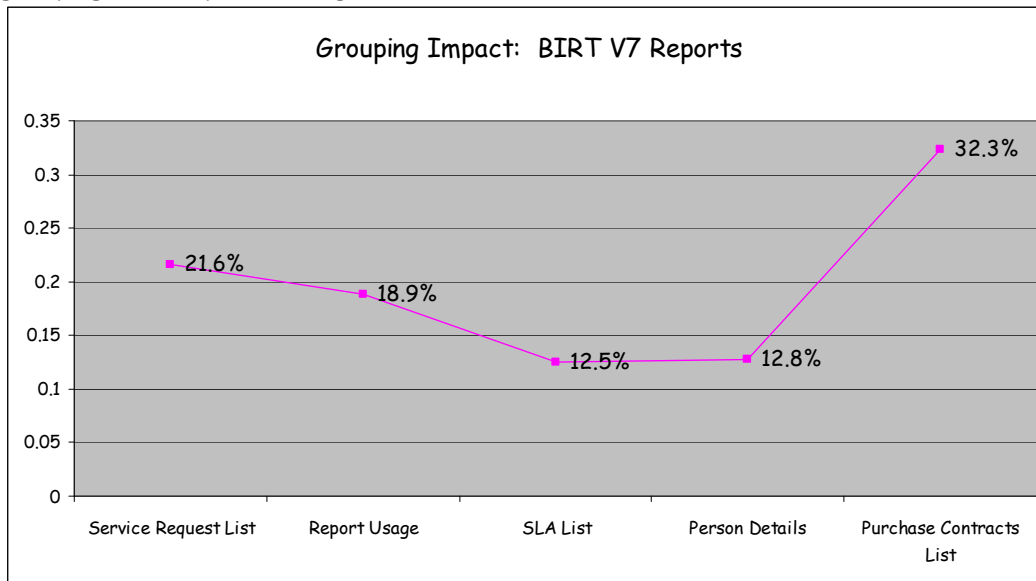
Asset	Description	Location	Bin	Parent	Type	Status	Priority	Rotating Item	Linear Asset?	Install Date	Mfr.
1008	Windows XP Operating System	OFF301		A7800		OPERATING		WINXP	N		
1009	Laser printer (local)	OFF301		A7800		OPERATING		HPLASER	N		
11210	Circulation Fan- Centrifugal/ 20/000 CFM	BR210		11200	FACILITIE S	OPERATING	4		N	5/31/94	TRN
11211	Motor Starter- Size 2/440v/3ph/60cy	BR210		11210		OPERATING	3		N	5/31/94	WES
11240	Circulation Fan- Centrifugal/ 20/000 CFM	BR240		11200	FACILITIE S	OPERATING	5		N	5/31/94	TRN
11250	Circulation Fan- Centrifugal/ 20/000 CFM	BR200		11200	FACILITIE S	OPERATING	4		N	6/2/98	TRN
11340	Motor Starter- Size 4/NEMA 12/440v/3ph/60hz	BR300		11300		NOT READY	4		N	5/31/94	WES

To quantify the performance impact of grouping in reports, five different reports were analyzed. These reports included Service Request List, Report Usage, SLA List, Person Details and Purchase Contract List. Two of these reports included two groupings, Service Request List and Report Usage, while the remaining three had one grouping.

Each report was first executed 10 times. Then, the groupings from each report were removed. The reports with no groupings were then executed 10 more times.

The results below show that on average, reports with no groupings ran 14% faster than those same reports that contained groupings.

The results below highlight the impact of groupings on reports. Although this is not as large an impact as graphs, it is an important factor to consider when specifying one or multiple groupings in a report's design.



## Report Controls

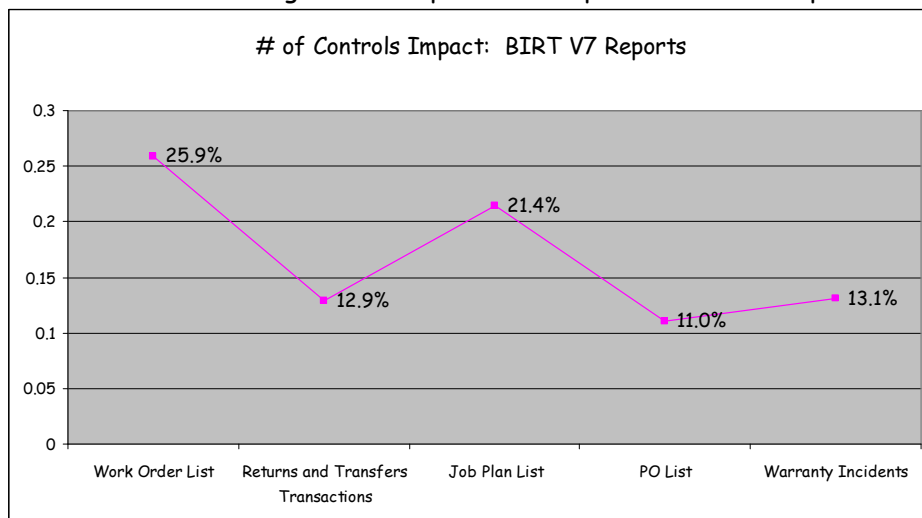
Every report contains multiple controls. These controls could be a database field, calculated value, graphic image or stationary line used to separate data results. The more controls designed in a report, the denser the report output becomes. And the denser a report, the more time it takes to execute. The screen shot below shows a dense report with 18 field controls in a single data row.

Date	Qty	From Location	To Location	From Lot	To Lot	From Bin	To Bin	Condition Code	Condition Code	From	Order Unit	Unit Cost	Line Cost	Actual Cost	Loaded Cost	PO	PO Line	Packing Slip	
12/2 4:38:00 PM	10	CENTRAL	CENTRAL	CH-4	N-3-8	A-4-9	G-5-3	A-4-9		EACH	129.00	1350.00	135.00	1417.50	A3250	2		52661	
1/19 1:33:00 PM	1	CENTRAL	CENTRAL			A-4-9	G-5-3			EACH	129.00	129.00	129.00						
11/2 1:00:00 PM	1	CENTRAL				A-4-9				EACH	139.00	139.00	139.00	0.00		5273	1	52661	

To quantify the impact of the number of controls used in reports, five different reports were analyzed. These reports included Work Order Tracking List, Returns and Transfers Transactions, Job Plan List, Purchase Order List and Warranty Incidents. Each report was executed 10 times.

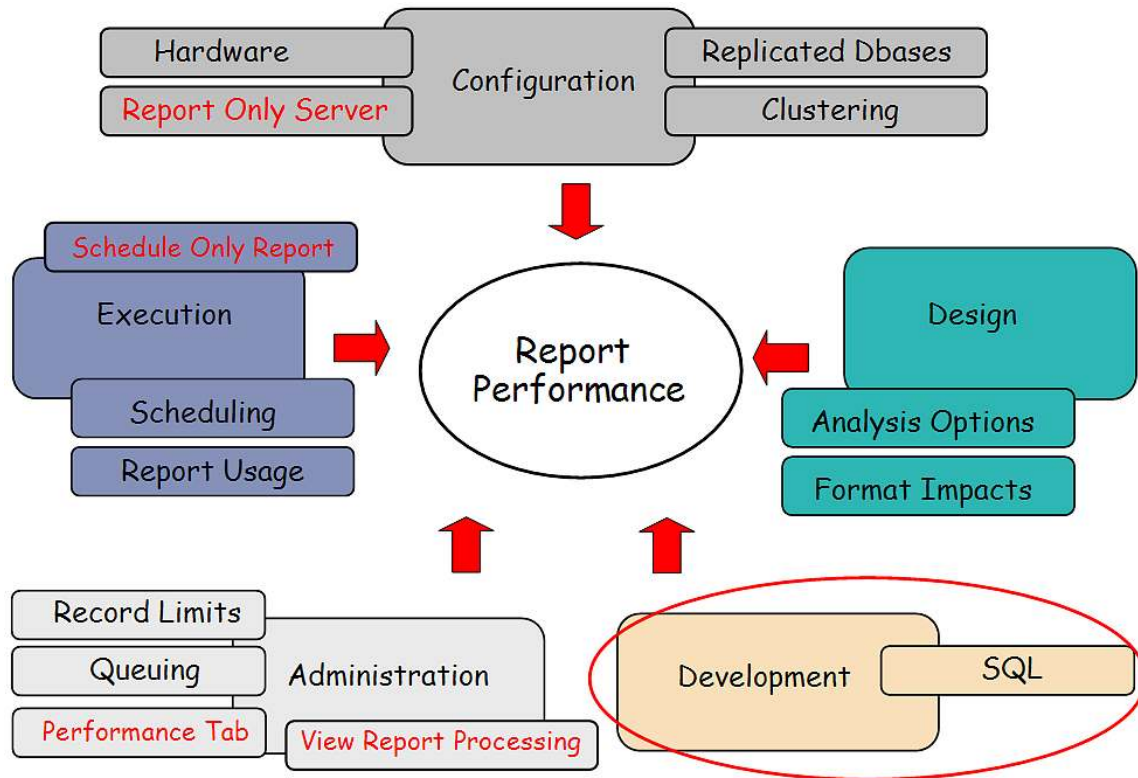
Then, the number of database field controls used in each report was reduced by 50%. For example in the Work Order Tracking List Report, the number of fields shown in the details section was reduced from 12 to 6. (Note: This did not reduce the number of overall controls used in the report by 50% - as many controls remain constant including the Report Titles, graphic images, page counts and date formats.) Depending on the individual report, the number of field controls was reduced from 9 to 4 fields. Once the database field controls were reduced by 50%, the 5 reports were then executed 10 times.

The results below show that on average, reports with reduced field controls ran 14% faster than those same reports that contained additional field controls. These results indicate that the number of fields designed in a report is an important factor of performance.



## Developing

In this third segment, the importance of optimizing the Report's sql will be discussed.



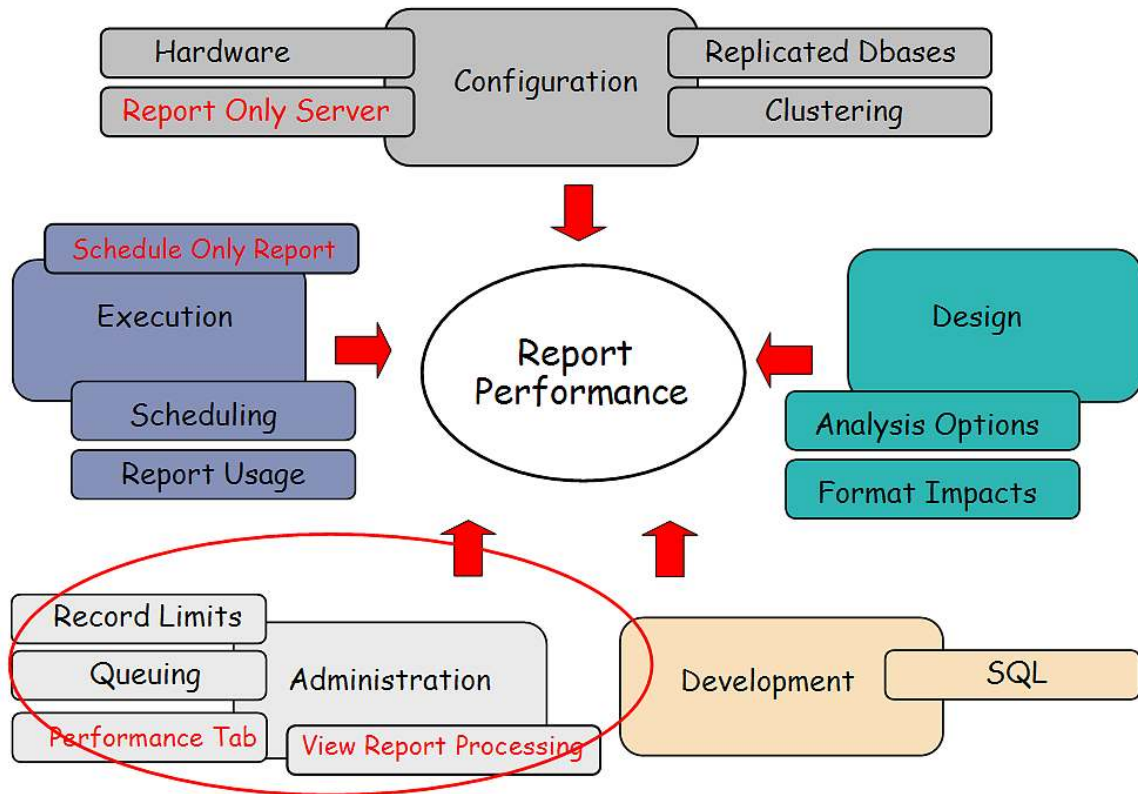
Each V7 Report contains SQL (Structured Query Language) which is used to query the database. The SQL details which database tables and fields need to be accessed. When more than one database table is accessed, joins are used to define the relationships between the tables.

The manner in which the data is retrieved from the database has a huge impact on performance. If a report retrieves data from multiple tables and the joins are incorrectly defined, the performance will be adversely impacted.

The BIRT Designer does not include any tools to either validate the report's sql statement or optimize it. Therefore, it is critical that the developer make sure that each of the report's sql is properly and optimally defined. Depending on the type of database used, the developer can use a variety of tools to do this. Additionally, the developer can verify the joins used in the report by evaluating the various maxrelationships. (These are defined in the V7 Database in the MAXRELATIONSHIPS Table.)

## Administration

Multiple actions can be performed in the Administering Process of the Reporting Cycle which both monitor performance and can be used to enhance performance, including (1) Record Limits (2) Queuing (3) Performance Tab and (4) View Report Processing.



## 1. Administration: Record Limits

Often times, users will inadvertently execute reports against large or very large record sets within an application. This can cause negative performance impacts on both the application and database server.

An example of this is in the Work Order Tracking application. In this case, a user only wants to print out the 10 new work orders which were input during the night shift. However, the user forgot to filter their query, and instead of submitting a Work Order Details Report Request for the only the 10 new work orders, they input a Work Order Details Report Request for all work orders, which totaled over 2500 records. And because the Work Order Details report is a very complex report, this request will tie up the application and database server for some time.

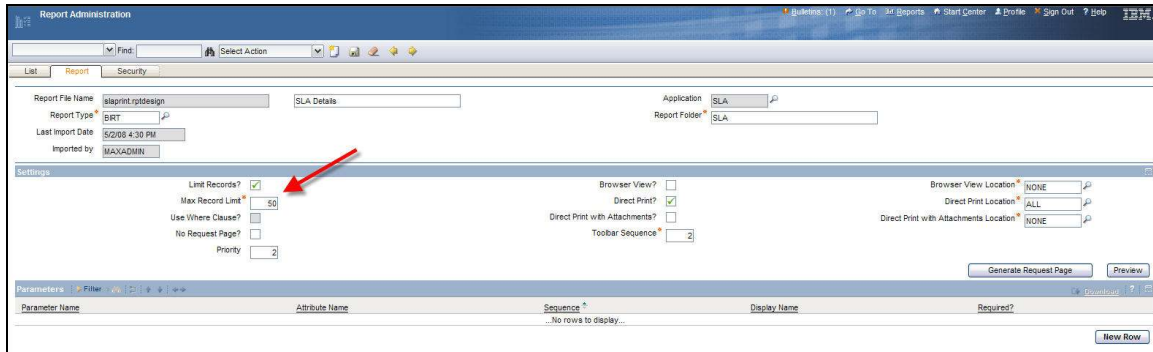
To prevent incidents like this, and from users running large queries, record limits can be enabled on V7 Reports. The Record Limit functionality enables you to define exactly the maximum number of records a report can execute against. This functionality is only enabled for reports that use the current/selected record set.

To enable this functionality, a count of the user's current record selection is determined from the application the user is in. This value is then checked against any record limits set for the selected report.

- If the current record set is greater than the record limit, a message is displayed to the user to reduce his query.
- If the current record set is less than the record limit, the report is executed.

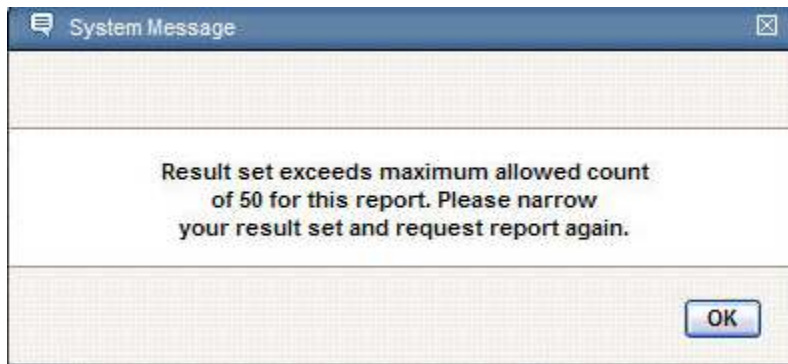
To set up Record Limits, access the Report Admin App.

1. Enable the Limit Record Flag
2. Input a value for the Max Record Limit. The value must be greater than 1, but there is no limit on the maximum number that can be input.
3. Save the record, and regenerate the XML.



The screenshot shows the 'Report Administration' interface. The 'Settings' section is expanded, and the 'Limit Records?' checkbox is checked. A red arrow points to this checkbox. The 'Max Record Limit' is set to 50. Other settings include 'Browser View?' (unchecked), 'Direct Print?' (checked), 'Priority' (2), and various location settings.

4. Then, if the user tries to execute a report from the app against a record set of 100 records - and the report has a limit of 50 - the error message below will display.



Note:

Only reports using Current/Selected parameters can take advantage of this feature. More information on the Business Rules for Record Limits is contained in the 'V7 or V75 Reporting Features Guide'. Information on accessing this guide can be found at this url:

<https://www.ibm.com/developerworks/wikis/display/maximo/Report+Reference+Materials>



## 2. Administration: Queuing

When numerous reports are executing in the report engine at once, performance can be adversely affected. This performance impact can be seen within the various applications, or in the amount of time it takes to process these reports.

Because of potential performance concerns, functionality has been developed which will set limits on the number of jobs that can be executed at any given time. This will distribute the load on the server from reports, and reduce the amount of wait time for report users.

Report Queuing will be used to handle the load of report jobs. The Queue Manager oversees the report queuing process. It is responsible for managing the scheduled jobs that enter the queue and the workers who process the scheduled jobs.

The report queuing process is made up of two components: scheduled report jobs and immediate reports jobs. Scheduled report jobs are controlled by the Queue Manager. The Queue Manager and worker threads do not apply to immediate report jobs. The Queue Manager and worker threads apply *ONLY* to scheduled report jobs.

To limit the number of report jobs that can be executed on a server at any time, the `maxconcurrentrun` property is used. By default this value is set to 5.

The Report Queuing process is best described with the diagram below, using the default value of 5 for the `maxconcurrentrun` property value.

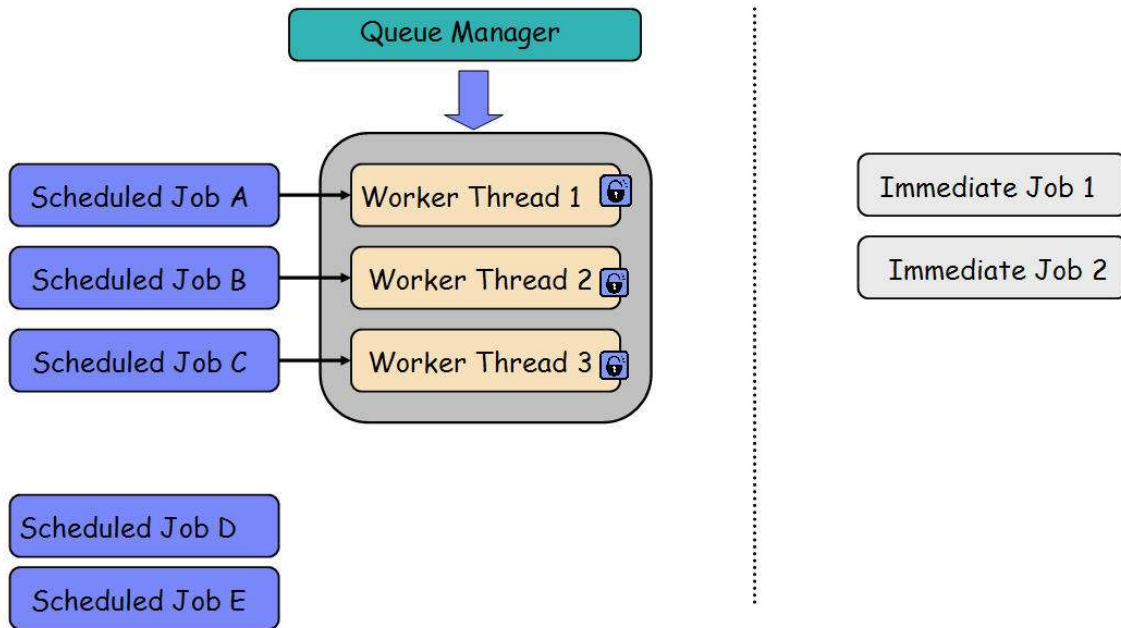
In this diagram, 2 immediate jobs are being processed, along with 3 scheduled report jobs. These 3 scheduled report jobs (A, B and C) are being controlled by the Queue Manager, and are being processed by 3 worker threads. The 2 immediate jobs are being processed by immediate job threads.

This is the maximum number of report jobs that can be processed at this time. This is determined by the `maxconcurrentrun` property value

**# Actively Running Scheduled Jobs + # Immediate Jobs  $\leq$  maxconcurrentrun**

In the diagram below, this is:

3 (Actively Running Scheduled Jobs) + 2 (Immediate Jobs)  $\leq$  5 (maxconcurrentrun property)



Additionally:

(1) When it is time for a scheduled report to be run, the scheduled report job enters the Queue.


Scheduled Jobs D and E have entered the reporting queue, but they can not be processed because there is no availability. (No available worker threads.)

(2) The Queue Manager checks the queue for scheduled jobs using the property setting, `queueidletimeseconds`.

- i. If jobs are present, then one job at a time is given to an available worker thread.
- ii. Scheduled Jobs are submitted to the Queue Manager based on
  1. Ascending Order of Report priorities from Levels 1 to 5
  2. Ascending Order of Scheduled Job Run Time
- iii. If there are no available worker threads, then the scheduled job will remain in the queue until a worker thread is available

Therefore, scheduled Jobs D and E will remain in the reporting queue until either one of the immediate or one of the scheduled report jobs completes.

This is shown in the updated diagram below where scheduled Job A has completed. The Queue Manager then moves scheduled Job D to worker thread 1, locks it and begins processing that D's report job request.



#### Notes

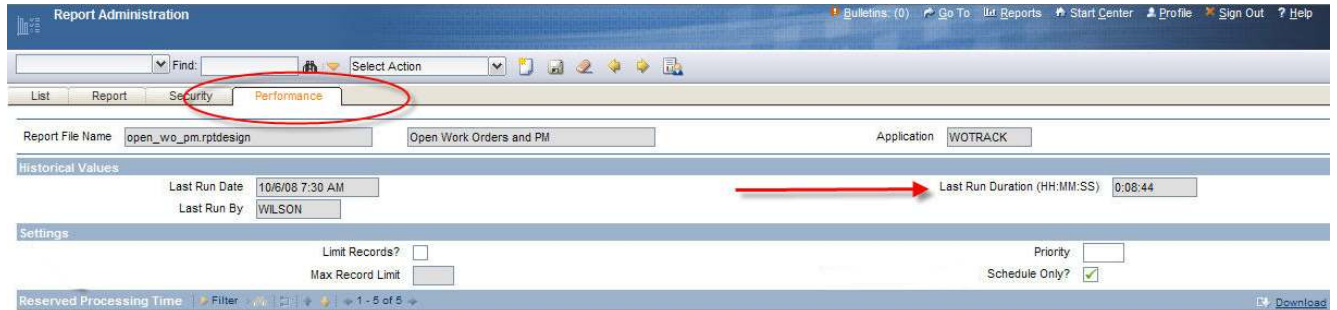
1. More information on the Property Files, Cron Tasks and Business Rules for Report Queuing is contained in the V7 or V75 Reporting Features Guide. Details on accessing these guides can be found at this url:

<https://www.ibm.com/developerworks/wikis/display/maximo/Report+Reference+Materials>

2. Also, for Report Only Server Configurations, please review the 'Configuration Recommendations and Notes' Section in this guide for maxconcurrentrun settings for this configuration option.

### 3. Administration: Performance Tab

Starting in Version 7.1.1.4, a tab is available to enable Report administrators to quickly monitor performance on individual reports.



The first section of the Performance Tab is historical values. This contains value information from the last time the report was executed. In the example below, it shows that the 'Open Work Orders and PM' Report registered in the Work Order Tracking application was last executed by user Wilson on 10/6/08 at 7:30am. The report executed in 8 minutes and 44 seconds.

This information is held in the REPORTUSAGELOG table and can also be displayed by executing the Report Usage Report.

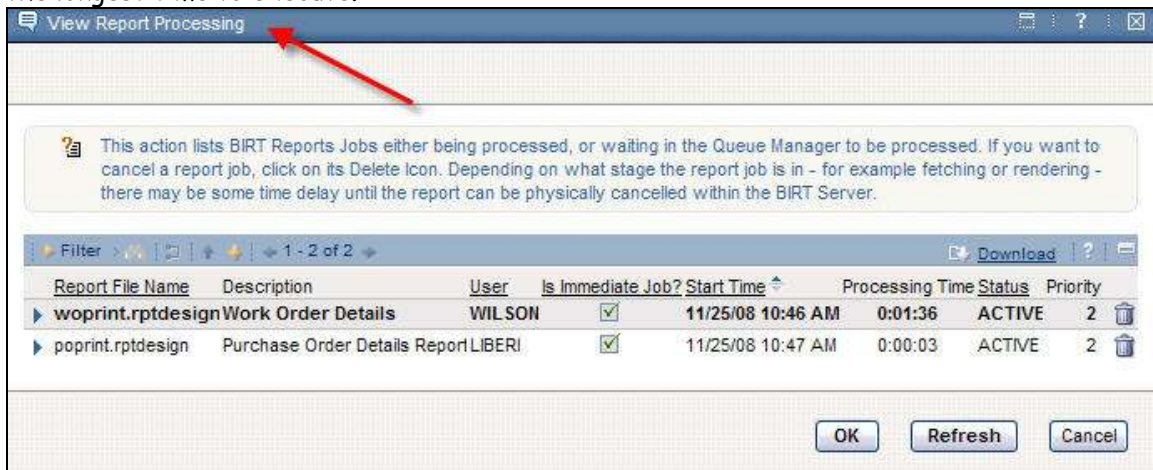
Enabling the display of these values quickly gives report administrators a data point on how long reports are taking to execute. This is valuable information in a development-type environment as administrators can identify long-executing reports, and enable features like 'Schedule Only' or 'Record Limits' to minimize their performance impacts.

The second section of the Performance Tab is Settings. This contains the settings that directly impact performance. These settings are also displayed on the Report Tab, but are also enabled here for ease of use. For example, if the administrator sees a long executing report, he can quickly set its 'Limit Records?' Flag here on this tab.

The third and final section of the Performance Tab is Reserved Processing. This functionality is enabled for 'Schedule Only' Reports. More details on this can be found at the end of this document.

## 4. Administration: View Report Processing

From within the Report Administration application, an action is available for the report administrator which displays reports being executed. Specifically, it shows all reports that are either being processed by the report engine, or are in the queue waiting to be processed. The records are displayed based on ascending order of Start Time. This display of records enables the administrator to focus on those report jobs that are taking the longest time to execute.



If the administrator clicks on the job details, other critical information including server and host name are displayed. The server and host name information are critical as the report job may not be processed on the same UI server that the end user submitted or scheduled the report.

For example, if a Cron Cluster is enabled for configured for report scheduling, (mxe.report.birt.disable.queuemanager property), then scheduled reports may execute on a different server.

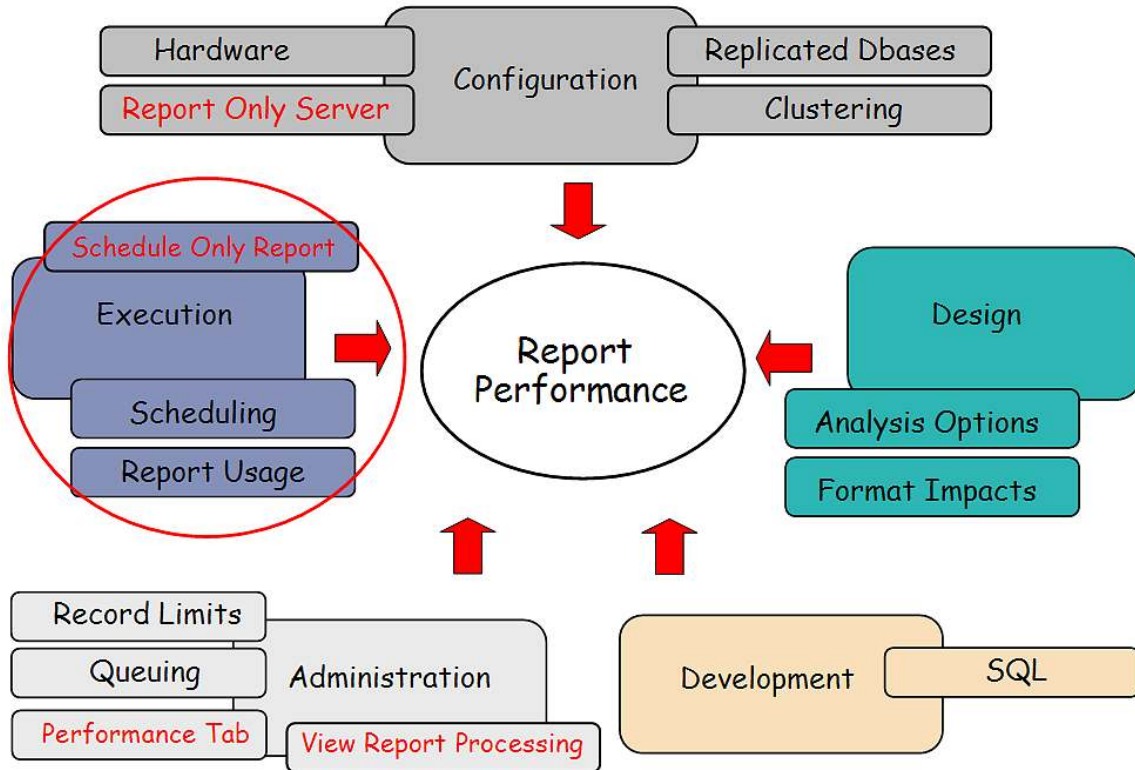
Or if you have configured a report only server via the mxe.report.birt.reportserver property, then immediate or scheduled reports may execute on a different server.

After viewing the report jobs that are processing, the Administrator has the ability to delete a report job by selecting the Garbage Can. Multiple events can occur when canceling an immediate or scheduled report job. Additionally, depending on the report process that is executing in the background determines how quickly the report can actually be cancelled.

Additional details on the business rules and the events that occur when canceling a Report Job can be found in the V7 Report Feature Guide.

## Execution

In the last component of the Report Process, Running Reports, the components of (1) Scheduling (2) Schedule Only Report and (3) Report Usage will be discussed.



## 1. Execution - Scheduling

In V7, Reports can be executed a variety of ways, including

- (1) run a report right now
- (2) run a report at a selected date/time in the future
- (3) run a report on a recurring basis

Running a report right now is called an immediate job, and passes a report job request to the BIRT Engine immediately. The other two options - running a report at a single point in the future - or on a future recurring basis - involve scheduling. Report scheduling is enabled through a Cron Task.

Request Page

Fill in the fields in the Parameters section below and click Submit to run the report. If no parameters are displayed, the report will execute against the current/selected record set. Optionally, fill out the Schedule and Email Sections to schedule report run times and receive the scheduled report via email.

Parameters

Schedule

Immediate

At this Time

Recurring

Email

To

Subject

Comments

Submit Cancel

Report Scheduling can be used as a performance enhancer by:

- A. Off-loading the report request to execute during Off-Peak Hours. Off-Peak Hours are defined as those hours in your individual business environment where the number of users actively using the system is minimized.

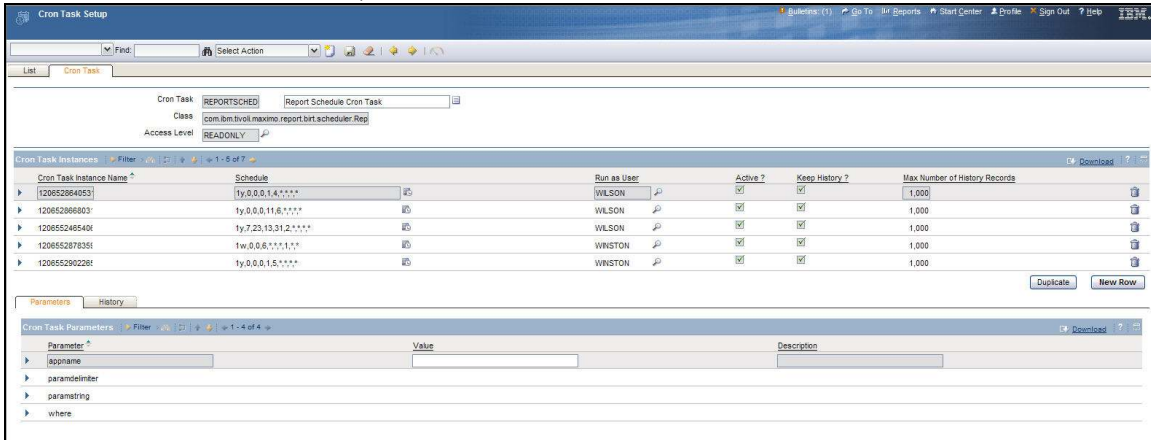
For example, scheduling a complex, analysis report like the Maintenance Cost Rollup to execute every Sunday at 4:00 AM may minimize the application and database performance impact rather than have the same report execute every Tuesday at 1:00 PM.

*Note: The Schedule Report Only - Reserved Processing Times functionality can be used to enforce the use of these off-peak hours.*

- B. Enabling the use of a Clustered cron server. By scheduling reports, the Clustered cron server can be used to execute the scheduled report versus the UI Server which is optimized to handle transactional requests and reports.

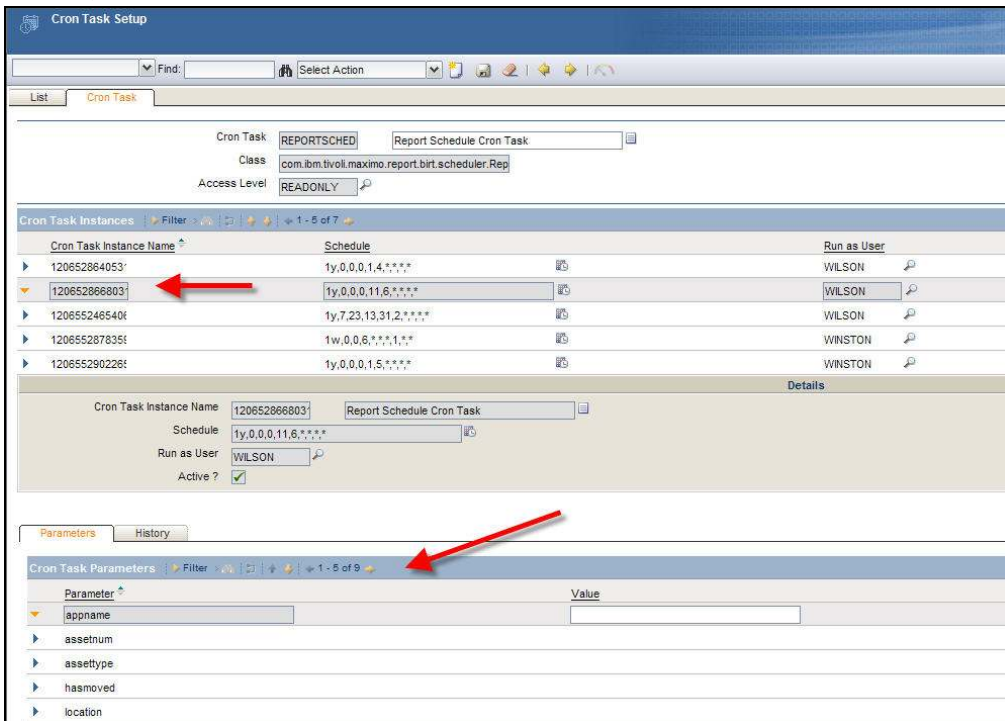


Executing the scheduling of reports is enabled thru a Cron Task Scheduler. The cron task is called REPORTSCHEDULE, and it references the class file that initiates it.



A separate cron task instance is created for each scheduled report. Therefore, the number of cron task instances is very dynamic as it is dependant on the number of scheduled jobs at any point in time.

The REPORTSCHEDULE cron task is read only, and is not intended to be modified or used to manage report schedules. Both the instance name and schedule values are auto-assigned values from the cron task.



Note: More information on Report Scheduling, including the Report Schedule Cron Task, the User's Scheduling Status and the Administration's View Schedule Report Action are contained in the V7 or V75 Report Feature Guides. Information on accessing these can be found at this url:

<https://www.ibm.com/developerworks/wikis/display/maximo/Report+Reference+Materials>




## 2. Execution: Schedule Only Reports

Beginning with Version 7.1.1.4, you can configure complex, batch reports to only be executed via report scheduling. This can prevent users from running large, complex reports during peak use of the application, database and report server, which can negatively impact overall system performance.

Before this functionality is described, let's review how this could be applied. In your environment, there is likely a wide range of complexity in reports. Some reports are very simple, like list reports or single grouped reports, whereas others are very complex due to the number of subreports they encompass or the hierarchy levels they span through. This Report complexity is defined by the processing the report has to do - not by the number of records it displays. So, for example, a 50 page list report could execute 10 times faster than a complex 2 page report due to the processing defined within the report's design.

The pie chart below represents a sample V7 Report Set. Standard, transactional reports are displayed in blue, with complex reports in red, and the most complex reports in yellow. These complex reports are time consuming to design and develop, and are continually analyzed for performance efficiencies. These are often referred to as batch reports, as they traditionally execute overnight due to their large processing requirements.



Because of the processing load complex reports have on a system, they can now be configured to only be executed via report schedules. Additionally, the very complex reports can be configured to only execute via schedules at specific times of the day. This will minimize the impact these complex reports have on overall system performance.

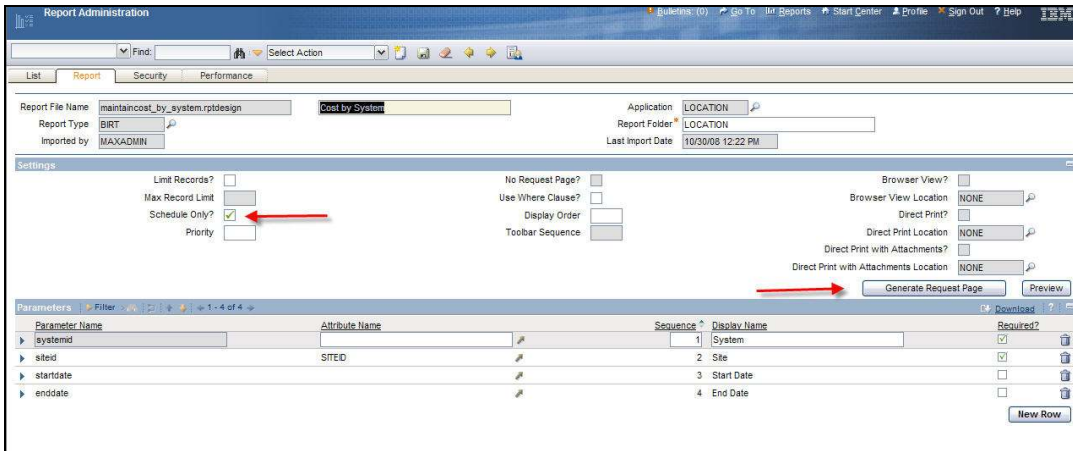
The next few pages describe how to enable this functionality.

To enable the 'Schedule Only' functionality for a complex report, access the Report Administration application and locate the report. The example below will use the Cost by System report.

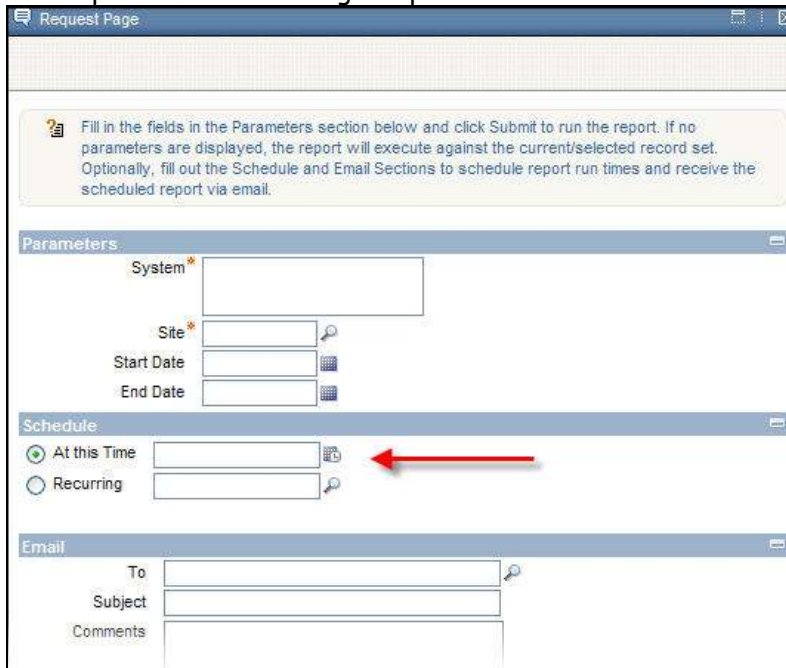
First, click Preview from the list tab to see what the report's request page looks like. Notice that three Execution Options display - Immediate, Schedule At this Time, or Schedule Recurring. These are the default run report options.

The screenshot shows a web application window titled "Request Page". At the top, there is a help icon and a text box with instructions: "Fill in the fields in the Parameters section below and click Submit to run the report. If no parameters are displayed, the report will execute against the current/selected record set. Optionally, fill out the Schedule and Email Sections to schedule report run times and receive the scheduled report via email." Below this are three main sections: "Parameters", "Schedule", and "Email". The "Parameters" section contains fields for "System\*", "Site\*", "Start Date", and "End Date". The "Schedule" section has three radio buttons: "Immediate" (which is selected), "At this Time", and "Recurring". A red arrow points to the "At this Time" radio button. The "Email" section has fields for "To", "Subject", and "Comments".

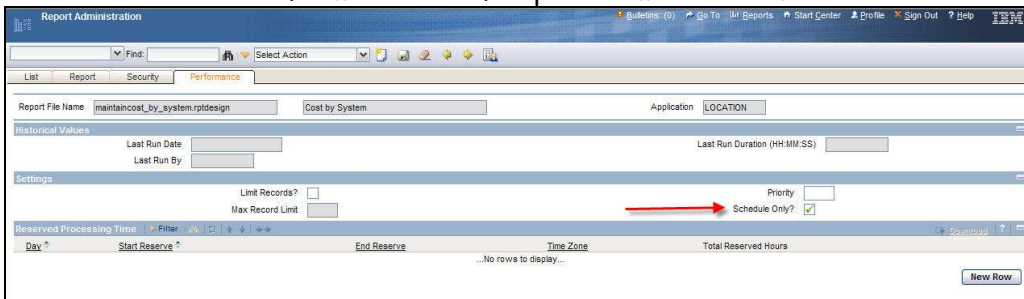
Now, enable this report for 'Schedule Only'. Close this window and on the Report Tab, enable the value for 'Schedule Only'. This value is not enabled by default.



Save the change, and generate the Request Page XML. Click Preview, and an updated request page displays. Now, the Immediate report execution option is no longer available. Only the two options for scheduling are presented to the user.



Finally, please note that the 'Schedule Only' functionality for a report can also be viewed and enabled on the Performance Tab of Report Administration.



## Notes on Functionality

1. If the Schedule Only field is enabled, the following fields can not be enabled:

No Request Page?

Browser View?

Direct Print?

Direct Print with Attachments?

This is because each of those fields relies on the request page NOT being displayed to the end user. In the case of scheduled only reports, the request page needs to be displayed so the schedule inputs can be made - therefore, these options are not available.

### 3. Execution - Schedule Only Reports - Reserved Processing Time

In addition to enabling Schedule Only functionality for very complex, batch reports, you can also configure what times are available for the report scheduling. This will prevent users from scheduling report jobs of complex reports at the peak times of server use. For example, you may not want your users to run very complex batch reports on Monday morning, when many users are trying to execute their daily or weekly transactional reports.

Reserved Processing time will detail the busy days and time within a week that server processing time is reserved for immediate report jobs and other maximo application functionality. This functionality will be based on days of the week. It is not based on calendar days, example January 1, due to fluctuations with calendars and holidays.

Reserved Processing Time can be set on any day of the week (Monday thru Sunday) and for any amount of time less than 24 hours. On the Performance Tab in Report Administration, the administrator will click New Row and specify the day of the week, and the times during that day that report processing is NOT available for that report job. This functionality is enabled by individual report because of the range of report complexity.

When the administrator clicks new row, multiple fields will display, including Day, Start Reserve, End Reserve and Total Reserved Time (Hrs).

The screenshot shows the 'Report Administration' interface, specifically the 'Performance' tab. The report file name is 'maintaincost\_by\_system.rptdesign' and the application is 'LOCATION'. The 'Reserved Processing Time' section is expanded, showing a table with columns for Day, Start Reserve, End Reserve, Time Zone, and Total Reserved Hours. The table lists reservations for Monday through Friday. A 'Details' section is visible below the table, showing the configuration for Monday: Start Reserve at 7:00 AM, End Reserve at 6:00 PM, and Total Reserved Hours of 11:00. The 'Schedule Only?' checkbox is checked.

Day	Start Reserve	End Reserve	Time Zone	Total Reserved Hours
MONDAY	7:00 AM	6:00 PM	America/Chicag	11:00
TUESDAY	7:00 AM	6:00 PM	America/Chicag	11:00
WEDNESDAY	7:00 AM	6:00 PM	America/Chicag	11:00
THURSDAY	7:00 AM	6:00 PM	America/Chicag	11:00
FRIDAY	5:00 AM	2:00 PM	America/Chicag	9:00

Details for Monday:  
 Day: MONDAY  
 Start Reserve: 7:00 AM  
 End Reserve: 6:00 PM  
 Time Zone: America/Chicag  
 Total Reserved Hours: 11:00

### What the Schedule Only Report Looks like to End User

If the 'Schedule Only' field is enabled for a report, the report can not be run immediately. The user will only have the choice to run this report either at a single point in the future, or on a recurring basis.

To highlight this functionality to the user

- A. The Immediate Runtime Option is hidden under the Scheduling Section
- B. The 'At this Time' Schedule Option will become the default.

The Schedule availability shown on the report's request page will show the opposite of what is shown in Report Administration. In Report Administration, the administrator is concerned with reserving time that the report should NOT be enabled to execute. Hence, it is called reserved processing time. However, for a report user, he wants to know when he can execute a report. He doesn't want to know when he can't execute a report - hence, the values will be shown as Schedule Availability.

The Report Availability also displays the schedule time availability time in the Time Zone in which it was SET. This may not match the time zone of the user.

Note: This functionality is expected to be enabled on complex reports executed by fairly technical users, who will understand time zones and their conversions.

The screenshot shows a web application window titled "Request Page". It contains several sections:

- Parameters:** Includes fields for "System" (a large yellow box), "Site" (with a magnifying glass icon), "Start Date", and "End Date". A red arrow points to the "Site" field.
- Schedule Availability:** A table with columns "Day", "Start Time", and "End Time". It lists days from SUNDAY to FRIDAY with their respective start and end times. Above the table are "Filter" and "Download" buttons, and a page indicator "1 - 10 of 12".
- Time Zone:** A section with the text "Available Time Shown in (GMT - 6 DST/Y) Central Standard Time."
- Schedule:** Includes radio buttons for "At this Time" (selected) and "Recurring", each followed by a text input field and a magnifying glass icon. A red arrow points to the "Recurring" radio button.

**Notes:**

1. To enable the Schedule Availability Section, the database table, REPORTPROCAVAIL, will be used. This table will essentially hold the opposite of the REPORTPROCRESERVE, which holds the report's reserved processing time. REPORTPROCAVAIL is required so end users can filter/search on Report Processing Availability from the Report's Request Page.

2. Due to potential conflicts with varying time zones - and the complexity involved with Daylight Savings time being utilized in some time zones, but not others - the administrator will only be able to reserve processing time for each report using the same time zone.

This means that if he reserves a processing time of 7 AM to 4 PM EST for the Asset Cost Rollup Report, any other reserved processing time made for Asset Cost Rollup must be made in the EST Time Zone.

3. The Schedule Only? Flag must be enabled for an administrator to set Reserved Processing Time.

4. Due to the wide range of unique client environments and requirements, no out of the box reports delivered have this functionality enabled. It is meant to be a configurable option that you can enable for their individual business environments.

5. After a Schedule Only report has been scheduled, its available processing time will not show on the Scheduling Status Tab. The report will show its next run date/time or recurring option in read only format only. If the report schedule needs to be modified, it is recommended that the original report schedule be deleted, and a new one input.

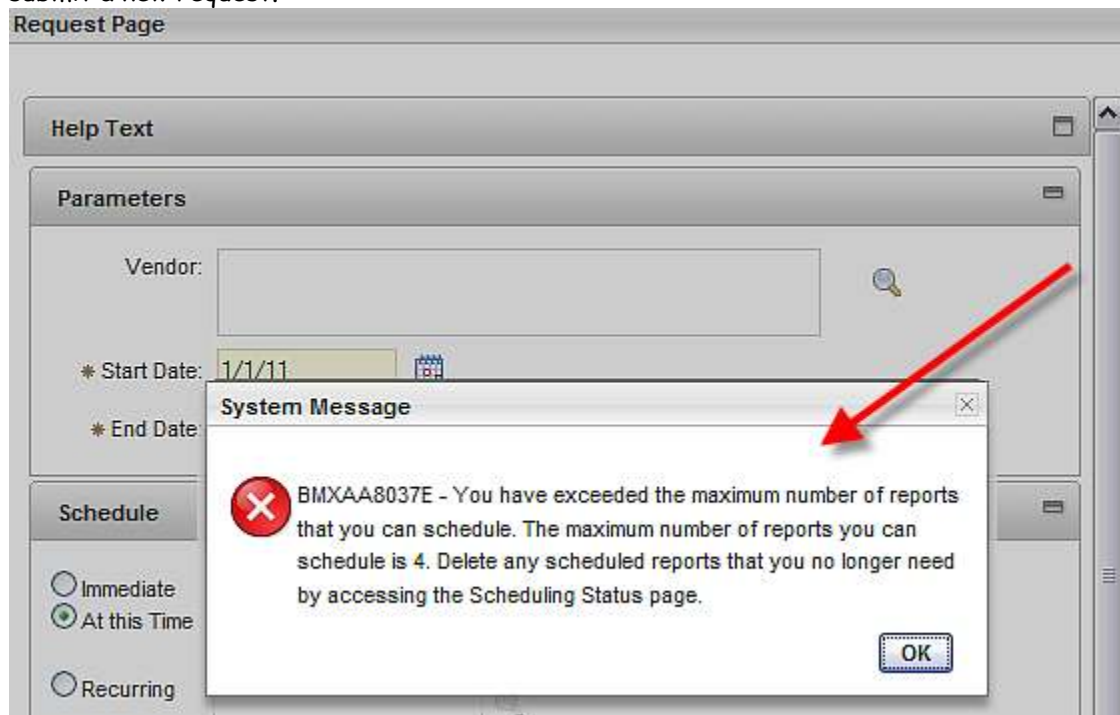


## Report Schedule Limits

You can limit the number of report schedule requests your users submit thru the report schedule limit functionality. By limiting them to a maximum number of scheduled jobs, users will only submit schedule requests for reports they need, and delete requests for obsolete report jobs.

Users may submit recurring schedule requests for reports that they think they will always need - but once the need for the report output goes away - they may forget to delete the scheduled report request. However, the report will still continue to execute. This leads to wasted, unnecessary resources - as the processing continues - while the consumption has stopped.

When scheduled report limits are enabled, and a user inputs a scheduled report request that exceeds his maximum allotted number, he will receive a message as shown below. He will then have to review his existing report requests, and delete one of them before he can submit a new request.



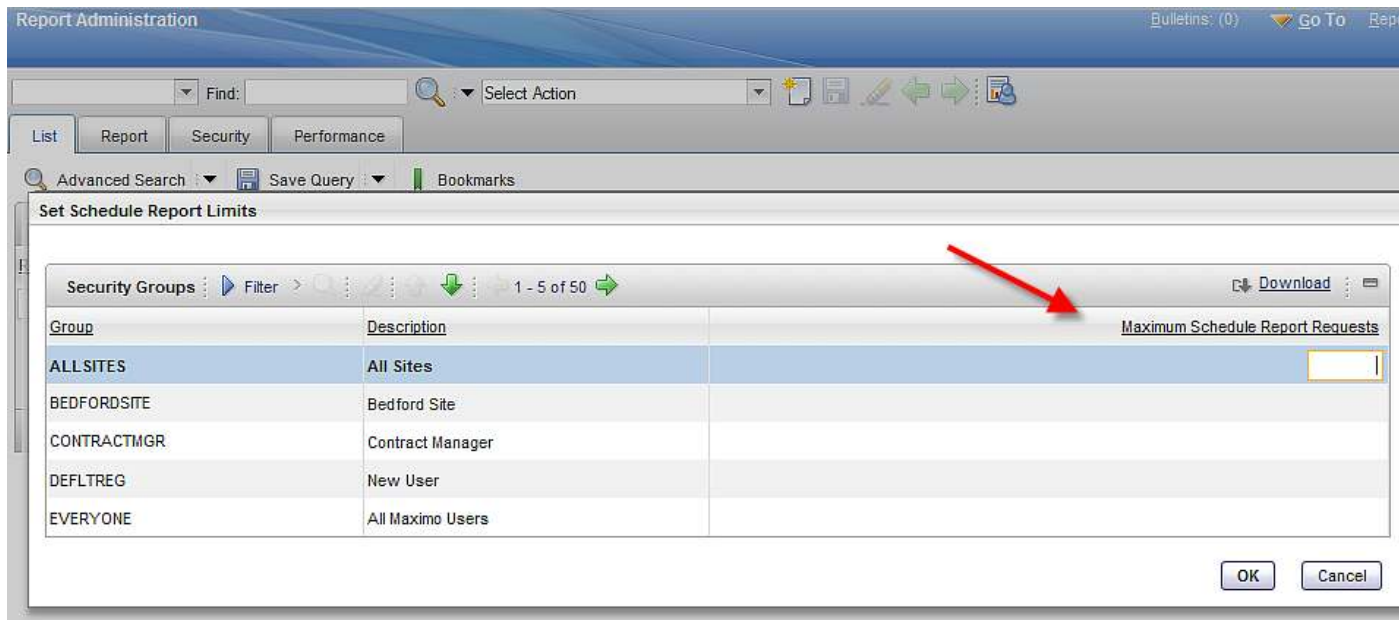
If a user is member of multiple security groups, the scheduled record limit for the user is the minimum value set for the various security groups.

This functionality has been introduced in the 7.5 release. Therefore, if you are upgrading from 7.1x to 7.5, your user's scheduled report requests will upgrade as is. If you then implement this functionality, it will not take affect until the user's next new scheduled report request.

## Set Report Schedule Limits

You can limit the number of report schedule requests your users submit thru the report schedule limit functionality. By limiting them to a maximum number of scheduled jobs, users will only submit schedule requests for reports they need, and delete requests for obsolete report jobs.

To configure this functionality, select the 'Set Schedule Report Limits' action. A dialog displays the security groups who have run report access. Input any number for each unique security group. If you input the number 0, the security group will not be able to input any scheduled report requests.



The example here shows that the Purchasing Security Group can enter a maximum of 4 scheduled report requests. This includes both enterprise and ad hoc scheduled report requests.

Report Administration Bullet

Find:  Select Action


List Report Security Performance

### Set Schedule Report Limits

Security Groups Filter >  1 - 1 of 1 Download

Group	Description	Maximum Schedule Report Requests
<input type="text" value="purchas"/>	<input type="text"/>	<input type="text"/>
PURCHASING	Purchasing	4


OK Cancel



## Managing Report Schedule Limits

You can manage the schedule limits you have set for individual security groups through a delivered report or KPI.

The Report Schedule Limits and Maintenance report details the number of security groups in the system. It then shows the value and percentage of those groups with run report access, and those with scheduled limits set. Administrators can use this report to quickly see the percentage of security groups that they closely managing.



The delivered KPI below can also be configured on the report administrator's start center to highlight those security groups with the schedule limit value set.

## 4. Execution - Report Usage

Functionality is available to monitor report use. This functionality enables you to monitor the usage and execution of BIRT reports. Report Usage Information can be used to answer the questions often raised by Development/IT including:

- Which reports take the longest to execute?'
- Does anyone really use this report? And if so, who are they?
- Are users taking advantage of the scheduling functionality? Or are they causing performance constraints by always running reports immediately?

The Report Usage report details the information collected from the Report Usage Table. This report is available from the Report Administration application. The report contains information on reports that are executed immediately, scheduled to execute at a future time, or are accessed via hyperlinks.

One of the most important fields in the Report Usage report is the Run Time Field. The data from this field comes from REPORTUSAGELOG.RUNTIME, which is stored in the database in milliseconds. The report converts the data from milliseconds to HH:MM:SS Format for ease of analysis.

Report Name: Application	Asset Details ASSET	File Name: asset_detail.rptdesign					
Start Date	End Date	Run Time (HH:MM:SS)	User	Application	Scheduled?	Success?	
3/27/08	7:28:47 AM	3/27/08 7:28:50 AM	00:00:03	WILSON	ASSET	N	Y

Report Name: Application	Asset List ASSET	File Name: asset.rptdesign					
Start Date	End Date	Run Time (HH:MM:SS)	User	Application	Scheduled?	Success?	
3/27/08	7:28:28 AM	3/27/08 7:28:40 AM	00:00:12	WILSON	ASSET	N	Y

Report Name: Application	Asset Measurement History ASSET	File Name: asset_measurement.rptdesign					
Start Date	End Date	Run Time (HH:MM:SS)	User	Application	Scheduled?	Success?	
3/27/08	7:30:53 AM	3/27/08 7:30:56 AM	00:00:03	WILSON	ASSET	N	Y

Report Name: Application	Vendor Contacts COMPANY	File Name: vendor_contacts.rptdesign					
Start Date	End Date	Run Time (HH:MM:SS)	User	Application	Scheduled?	Success?	
3/27/08	7:35:53 AM	3/27/08 7:35:54 AM	00:00:01	WILSON	COMPANY	N	Y

Query Based Reports (QBRs) are also captured in the Report Usage Table. Because these reports often may be executed only a single time, their corresponding run times will be displayed in a separate section, called 'Transient Reports' at the end of this Report.

Transient Reports							
	Start Date	End Date	Run Time (HH:MM:SS)	User	Application	File Name	Success?
	3/27/08 7:51:53 AM	3/27/08 7:52:00 AM	00:00:07	SNOLAN	WOTRACK	snolan1206616979250.rptdesign	Y
	3/27/08 7:31:50 AM	3/27/08 7:31:52 AM	00:00:02	WILSON	JOBPLAN	wilson1206617776921.rptdesign	Y
	3/27/08 7:31:17 AM	3/27/08 7:31:21 AM	00:00:04	WILSON	ASSET	wilson1206617744406.rptdesign	Y

## Additional References

The following lists additional references available at the time this document was prepared.

### **Report Wiki Site**

Dynamic location to access the latest information on reporting.

<https://www.ibm.com/developerworks/wikis/display/maximo/Reports>

### **Report Wiki Page - Listing all reporting documentation by release**

<https://www.ibm.com/developerworks/wikis/display/maximo/Report+Reference+Materials>

### **Report Wiki Page - Listing all Enhancements by Release**

<https://www.ibm.com/developerworks/wikis/display/maximo/Major+Report+Enhancements+by+Release+Details>

### **IBM's Support Site.**

<http://www-01.ibm.com/software/sysmgmt/products/support/IBMMaximoAssetManagement.html>

### **System Administration Guide**

System Configuration Chapter: Contains details on Advanced Server Configurations

<http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=/com.ibm.itmaxam.doc/welcome.htm>

## Revision History

Revision 12: June 2012 (1) Updated Execution portion on page 49 to include Version 7.5's capability of setting schedule report limits for users to manage scheduled report jobs (2) Added note on page 17 that LDAP is supported with BROS and Maximo if they both have the same domain name.

Revision 11: February 2012 (1) Added table of contents (2) Minor ui and text updates (3) Updated data analysis section (4) Updated reference materials

Revision 10: January 2011 (1) Added note that the BROS server does not support application server security on page 17.

Revision 9: December 2010 (1) Added note on page 10 to functionality exceptions to BROS server

Revision 8: October 2010 (1) Added http to url's used in examples in BROS section (2) Included additional notes on BROS on page 11

Revision 7: June 2010 (1) Updated diagrams on pages 11 and 12



## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

### **Trademarks**

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “[Copyright and trademark information](#)” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both