Platform™ Analytics™ Cube Descriptions

Platform

Contents

1

The Daily FLEX License Cube

The Daily FLEX License cube lets you explore historical daily data about FLEX license usage and other events contained in the FLEXnet Manager server. This cube is based on the Daily FLEX License datamart.

Note: Macrovision's FLEXnet Manager is a central management tool for all FLEXlm servers in an enterprise environment. In order to build and explore the data that makes up this cube, you require a FLEXnet Manager installation.

Contents
- "Dimensions" on page 6
- "Measures" on page 8
- "Cognos Cube Views" on page 9

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily FLEX License cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday

	◆ Friday
	◆ Saturday
Business Hours	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Project	◆ [User-defined Project]
License Server	◆ [License Server]
Vendor	◆ All vendors | [License vendor] | [License feature]
Version	◆ All versions | [License version]
LSF/non-LSF	◆ LSF
	◆ Non-LSF
Event	◆ DENIED
	◆ USED
Event Reason	◆ Licensed number of users already reached
	◆ Normal
	◆ UNKNOWN
User	◆ [All Users] | [User]
Host	◆ All Hosts | [Host]

Measures

The Daily FLEX License cube includes the following measures:

of checkouts

Number of license checkouts.

Max # of checkouts

Top number of license checkouts at one time for each job; measures the frequency of checkouts for a single application.

Total feature time (min)

Total license feature time per job, in minutes; combines all consumed license time.

Avg feature time (min)

Average license feature time per job, in minutes; averages the consumed license time.

of events

Number of events (for example, denial, used).

Cognos Cube Views

The Daily FLEX License cube does not include any Cognos cube views.

2

The Daily Hardware Cube

The Daily Hardware cube lets you explore the load on your cluster and ensure that it is configured correctly and fully utilized. It contains daily historical hardware data. It is based on the Host Properties and Daily Hardware datamarts. You can view data by year, quarter, month, week, day, or hour. You can explore hardware usage for all hosts, host models, and host types by day.

Contents
- "Dimensions" on page 12
- "Measures" on page 14
- "Cognos Cube Views" on page 15

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Hardware cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday

	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	The number of cores in each physical processor in the host.
	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	The maximum number of threads in each core in the host.
	◆ 1
	◆ 2

Measures

The Daily Hardware cube includes the following measures:

# of logins	Number of logins.
CPU (%)	Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.
Used mem (%)	Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).
Used swap (%)	Used swap space as a percentage.
Used tmp (%)	Used temp space as a percentage.
15s run queue	Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.
1min run queue	Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.
15min run queue	Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.
Paging (pgs io/s)	Paging activity, in pages in and out per second.
Disk IO (KB/second)	Average disk I/O, for all disks on a host, in kilobytes per second.
Inter. idle time (min)	Number of minutes since any interactive terminal activity.
Free mem (MB)	Available physical memory, in megabytes.
Free swap (MB)	Available swap space, in megabytes.
Free tmp (MB)	Available space in a temporary file system, in megabytes.
Max mem (MB)	Average total amount of maximum physical memory, in megabytes.
Max swap (MB)	Average total amount of maximum swap space, in megabytes.
Max tmp (MB)	Average total amount of maximum space in a temporary file system, in megabytes.
Used mem (MB)	Used physical memory, in megabytes.
Used swap (MB)	Used swap space, in megabytes.
CPU (%) used by jobs	Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.
Job used mem (MB)	Total of all physical memory used for jobs, in megabytes.
Job used swap (MB)	Total of all swap space used for jobs, in megabytes.

Cognos Cube Views

The Daily Hardware cube includes the following Cognos cube views:

- Quarterly performance by host type
- System profile by host type
- Which hosts need more disk space
- Which hosts need more memory

3

The Daily Hardware by Hostgroup Cube

The Daily Hardware by Hostgroup cube contains daily historical hardware data by host group. You can explore hardware usage for all host groups by day. This cube is based on the Daily Hardware by Hostgroup datamart. View data by year, quarter, month, week, day, or hour.

Note: This cube is not enabled by default due to the large of amount of data collected. To build this cube, check the appropriate box within the **Scheduled Tasks** window. See *Administering Platform Analytics* for more information.

Contents

Dimensions

Abbreviations

WTD Week to Date

MTD Month to Date

QTD Quarter to Date

YTD Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Hardware by Hostgroup cube includes the following dimensions and child categories:

Host Group Host group information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Group] | [Host]

Note: If a host belongs to multiple host groups, then the data is duplicated in each group.

Cluster Cluster | [Host]

Date The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Measures

The Daily Hardware by Hostgroup cube includes the following measures:

# of logins   Number of logins.

CPU (%)   Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

Used mem (%)   Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).

Used swap (%)   Used swap space as a percentage.

Used tmp (%)   Used temp space as a percentage.

15s run queue   Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

1min run queue   Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

15min run queue   Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Paging (pgs io/s)   Paging activity, in pages in and out per second.

Disk IO (KB/second)   Average disk I/O, for all disks on a host, in kilobytes per second.

Inter. idle time (min)   Number of minutes since any interactive terminal activity.

Free mem (MB)   Available physical memory, in megabytes.

Free swap (MB)   Available swap space, in megabytes.

Free tmp (MB)   Available space in a temporary file system, in megabytes.

Max mem (MB)   Average total amount of maximum physical memory, in megabytes.

Max swap (MB)   Average total amount of maximum swap space, in megabytes.

Max tmp (MB)   Average total amount of maximum space in a temporary file system, in megabytes.

Sampling count   The weight used to calculate average values of other measures. This measure is for internal use only.

Used mem (MB)   Used physical memory, in megabytes.

Used swap (MB)   Used swap space, in megabytes.

CPU (%) used by jobs   Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.

Job used mem (MB)   Total of all physical memory used for jobs, in megabytes.

Job used swap (MB)   Total of all swap space used for jobs, in megabytes.

Cognos Cube Views

The Daily Hardware by Hostgroup cube includes the following Cognos cube views:

- Quarterly performance by host type
- System profile by host type
- Which hosts need more disk space
- Which hosts need more memory

4

The Daily Hardware (with Boolean Resources) Cube

The Daily Hardware (with Boolean Resources) cube lets you explore the load on your cluster and ensure that it is configured correctly and fully utilized. It contains daily historical hardware data. It is based on the Host Properties and Daily Hardware datamarts. You can view data by year, quarter, month, week, day, or hour. You can explore hardware usage for all hosts, host models, and host types by day. For parallel jobs, the cube considers the first host in the execution host list as the execution host.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Hardware (with Boolean Resources) cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Boolean Resource	Host-based Boolean resources. When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled.
# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4

# of Cores per CPU	The number of cores in each physical processor in the host.
	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	The maximum number of threads in each core in the host.
	◆ 1
	◆ 2

Measures

When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled. The Daily Hardware (with Boolean Resources) cube includes the following measures:

# of logins Number of logins.

CPU (%) Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

CPU factor The relative CPU performance factor. The CPU factor is used to scale the CPU load value so that differences in CPU speeds are considered. The faster the CPU, the larger the CPU factor.

Used mem (%) Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).

Used swap (%) Used swap space as a percentage.

Used tmp (%) Used temp space as a percentage.

15s run queue Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

1min run queue Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

15min run queue Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Paging (pgs io/s) Paging activity, in pages in and out per second.

Disk IO (KB/second) Average disk I/O, for all disks on a host, in kilobytes per second.

Inter. idle time (min) Number of minutes since any interactive terminal activity.

Free mem (MB) Available physical memory, in megabytes.

Free swap (MB) Available swap space, in megabytes.

Free tmp (MB) Available space in a temporary file system, in megabytes.

Max mem (MB) Average total amount of maximum physical memory, in megabytes.

Max swap (MB) Average total amount of maximum swap space, in megabytes.

Max tmp (MB) Average total amount of maximum space in a temporary file system, in megabytes.

Used mem (MB) Used physical memory, in megabytes.

Used swap (MB) Used swap space, in megabytes.

CPU (%) used by jobs Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.

Job used mem (MB) Total of all physical memory used for jobs, in megabytes.

Job used swap (MB) Total of all swap space used for jobs, in megabytes.

Cognos Cube Views

The Daily Hardware (with Boolean Resources) cube includes the following Cognos cube views:

◆ Quarterly performance by host type
◆ System profile by host type
◆ Which hosts need more disk space
◆ Which hosts need more memory

5

The Daily Hardware (with Boolean Resouces and Core Utilization) Cube

The Daily Hardware (with Boolean Resources and Core Utilization) cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Daily Hardware (with Boolean Resources and Core Utilization) and Hourly Core Utilization datamarts. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, and host types by day. For parallel jobs, the cube considers the first host in the execution host list as the execution host.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware (with Core Utilization) cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Boolean Resource	Host-based Boolean resources. When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled.

# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	◆ 1
	◆ 2
Host–Core ID	Analytics-assigned ID that is unique to each core within each host.
	◆ [Host-Core ID]

Measures

When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled. The Hourly Hardware (with Core Utilization) cube includes the following measures:

# of logins Number of logins.

CPU ut (%) Average CPU utilization as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

Used mem (%) Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).

Used swap (%) Used swap space as a percentage.

Used tmp (%) Used temp space as a percentage.

15s run queue Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

1min run queue Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

15min run queue Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Paging (pgs io/s) Paging activity, in pages in and out per second.

Disk IO (KB/second) Average disk I/O, for all disks on a host, in kilobytes per second.

Inter. idle time (min) Number of minutes since any interactive terminal activity.

Free mem (MB) Available physical memory, in megabytes.

Free swap (MB) Available swap space, in megabytes.

Free tmp (MB) Available space in a temporary file system, in megabytes.

Max mem (MB) Average total amount of maximum physical memory, in megabytes.

Max swap (MB) Maximum available swap space, in megabytes.

Max tmp (MB) Maximum available space in a temporary file system, in megabytes.

Used mem (MB) Used physical memory, in megabytes.

Used swap (MB) Used swap space, in megabytes.

CPU (%) used by jobs Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.

Job used mem (MB) Total of all physical memory used for jobs, in megabytes.

Job used swap (MB) Total of all swap space used for jobs, in megabytes.

Core ut (%) Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.

# of hosts Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.

Cognos Cube Views

The Daily Hardware (with Core Utilization) cube includes the following Cognos cube views:

- Daily CPU utilization trends by host type
- Which hosts run too many jobs

6

The Daily Hardware (with Core Utilization) Cube

The Daily Hardware (with Core Utilization) cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Daily Hardware (with Core Utilization) and Hourly Core Utilization datamarts. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, and host types by day. For parallel jobs, the cube considers the first host in the execution host list as the execution host.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware (with Core Utilization) cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]

# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	◆ 1
	◆ 2
Host-Core ID	Analytics-assigned ID that is unique to each core within each host.
	◆ [Host-Core ID]

Measures

When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled. The Hourly Hardware (with Core Utilization) cube includes the following measures:

# of logins Number of logins.

CPU ut (%) Average CPU utilization as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

Used mem (%) Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).

Used swap (%) Used swap space as a percentage.

Used tmp (%) Used temp space as a percentage.

15s run queue Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

1min run queue Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

15min run queue Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Paging (pgs io/s) Paging activity, in pages in and out per second.

Disk IO (KB/second) Average disk I/O, for all disks on a host, in kilobytes per second.

Inter. idle time (min) Number of minutes since any interactive terminal activity.

Free mem (MB) Available physical memory, in megabytes.

Free swap (MB) Available swap space, in megabytes.

Free tmp (MB) Available space in a temporary file system, in megabytes.

Max mem (MB) Average total amount of maximum physical memory, in megabytes.

Max swap (MB) Maximum available swap space, in megabytes.

Max tmp (MB) Maximum available space in a temporary file system, in megabytes.

Used mem (MB) Used physical memory, in megabytes.

Used swap (MB) Used swap space, in megabytes.

CPU (%) used by jobs Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.

Job used mem (MB) Total of all physical memory used for jobs, in megabytes.

Job used swap (MB) Total of all swap space used for jobs, in megabytes.

Core ut (%) Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.

# of hosts Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.

Cognos Cube Views

The Daily Hardware (with Core Utilization) cube includes the following Cognos cube views:

- Daily CPU utilization trends by host type
- Which hosts run too many jobs

7

The Daily Host Availability Cube

The Daily Host Availability cube contains daily historical host down-time data. You can explore down-hours and down-times for all hosts, host models, and host types by day. View data by year, quarter, month, week, day, or hour. This cube is based on the Daily Host Availability and Host Properties datamarts.

Contents   ◆   "Dimensions" on page 42
       ◆   "Measures" on page 44
       ◆   "Cognos Cube Views" on page 45

Dimensions

Abbreviations

WTD Week to Date

MTD Month to Date

QTD Quarter to Date

YTD Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Host Availability cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]

Measures

The Daily Host Availability cube includes the following measures:

Down hours Number of hours that a host was down during a specified time period.

Down times Number of times that a host went down during a specified time period.

> **Note**: There may be instances when the host is already down when the agent begins collecting host status data. If the host stays down during the whole data collection period (for example, one hour), than the down time will be reported as "0" (meaning it has not *gone* down during that time period).

Down (%) Down hours divided by a period of time.

Cognos Cube Views

The Daily Host Availability cube does not include any Cognos cube views.

8

The Daily Host Status Cube

The Daily Host Status cube contains daily historical host status data. You can explore the how long a host remains in a certain state for all hosts, host models, and host types by day. View data by year, quarter, month, week, day, or hour. This cube is based on the Daily Host Status and Host Properties datamarts.

Contents

- "Dimensions" on page 48
- "Measures" on page 50
- "Cognos Cube Views" on page 51

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Host Status cube includes the following dimensions and child categories:

Status Status Name

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]

Measures

The Daily Host Status cube includes the following measure:

Duration (min) — Length of time (in minutes) the host remained in a certain state during a specified time period.

Cognos Cube Views

The Daily Host Status cube does not include any Cognos cube views.

9

The Daily Job Resource Consumption Cube

The Daily Job Resource Consumption cube lets you explore historical data about resources reserved and consumed by jobs. A job resource is the LSF resource used in the resource string when submitting LSF jobs with the `-R` option, specifically in the `rusage [ ]` section. It is based on data collected from the `lsb.acct` file, collected in the Daily Job Reserved Resource Usage datamart.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Job resource   A resource that can be reserved. Even if `decay` or `duration` is used in the resource string, the entire run time is still added to the job.

> **Note:** To track how a license is reserved, you must specify it in the `rusage` section.

Job resource consumption   Job resource consumption equals the total number of minutes that all resources are reserved.

Dimensions

The Daily Job Resource Consumption cube includes the following dimensions and child categories:

Resource   Resource | [Reserved Resource]

Cluster   Cluster | [Host]

Date   The date when the job finished running, based on the date and time according to the LSF master host.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆   [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆   [Yesterday]
◆   WTD | [Day]
◆   MTD | [Week] | [Day]
◆   QTD | [Month] | [Week] | [Day]
◆   YTD | |[Month] | [Week] | [Day]

♦ [Previous Week] | [Day]

♦ [Previous Month] | [Week] | [Day]

♦ [Previous Quarter] | [Month] | [Week] | [Day]

♦ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

♦ Sunday

♦ Monday

♦ Tuesday

♦ Wednesday

♦ Thursday

♦ Friday

♦ Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

♦ Core hours

♦ Non-core hours

Project   ♦ [User-defined project]

User   ♦ [All Users] | [User]

Measures

The Daily Job Resource Consumption cube includes the following measures:

Usage (minute) The total number of minutes reserved by all job resources. By default, this shows the amount of time in minutes that each resource is reserved. (Resource names are rows, and cluster names are columns.)

Cost Cost is calculated as **Usage (minute)** × **Resource Rate**. The Resource Rate is defined in the console under **Clusters > [*Cluster Name*] > Resource Costs > Job Cost Rate**. You can specify the job cost rate as well as the start date for applying this rate.

Cognos Cube Views

The Daily Job Resource Consumption cube includes the following Cognos cube view report:

◆ Weekly resource consumption trend

10

The Daily Job Slot Usage Cube

The Daily Job Slot Usage cube lets you explore historical daily job slot utilization and usage data according to group type. It is based on the Daily Job Slot Usage by Group datamart. You can explore job slot utilization and usage information by user, host, project, queue, job group, and job name.

Contents ◆ "Dimensions" on page 60
◆ "Measures" on page 62
◆ "Cognos Cube Views" on page 63

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Job Slot Usage cube includes the following dimensions and child categories:

Group Type
- By Cluster
- By Host
- By Jobgroup
- By Jobname
- By Project
- By Queue
- By User
- By Usergroup

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]

- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Measures

The Daily Job Slot Usage cube includes the following measures:

Max Slot Usage
: Top number of slots used at one time for each group type; measures the number of job slots used for a group type.

Min Slot Usage
: Minimum number of slots used at one time for each group type; measures the number of job slots used for a group type.

Avg Slot Usage
: Average number of slots used at one time for each group type; measures the number of job slots used for a group type.

Max Slot Utilization (%)
: The maximum slot utilization for each group type, expressed as a percentage.

Min Slot Utilization (%)
: The minimum slot utilization for each group type, expressed as a percentage.

Avg Slot Utilization (%)
: The average slot utilization for each group type, expressed as a percentage.

Cognos Cube Views

The Daily Job Slot Usage cube does not include any Cognos cube views.

11

The Daily Job Statistics Cube

The Daily Job Statistics cube lets you explore historical job data about jobs in the cluster in any state (running, pending, and suspended). It is based on the Daily Job Statistics datamart. You can explore the job statistics for all hosts, projects, queues, users, and clusters.

Contents

- "Dimensions" on page 66
- "Measures" on page 69
- "Cognos Cube Views" on page 70

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Notes about the Group, Job Type, and Job Status dimensions

In order to obtain meaningful data, you must specify child categories (i.e. filter) all of the following dimensions:

- Group
- Job Type
- Job Status

You can filter these dimensions from the Dimension line. For additional information, see "Filtering data" on page 67.

Dimensions

The Daily Job Statistics cube includes the following dimensions and child categories:

Group   To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 67.

- By Host
- By Project
- By Queue
- By User
- By Cluster

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]

◆ [Yesterday]
◆ WTD | [Day]
◆ MTD | [Week] | [Day]
◆ QTD | [Month] | [Week] | [Day]
◆ YTD | |[Month] | [Week] | [Day]
◆ [Previous Week] | [Day]
◆ [Previous Month] | [Week] | [Day]
◆ [Previous Quarter] | [Month] | [Week] | [Day]
◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours
◆ Non-core hours

Job Type   To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 67.

◆ Batch
◆ Parallel

Job Status   To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 67.

◆ Pending
◆ Running
◆ Suspended
◆ Waiting

Weekday   Based on the date and time according to the LSF master host.

◆ Sunday
◆ Monday
◆ Tuesday
◆ Wednesday
◆ Thursday
◆ Friday
◆ Saturday

Filtering data

In order to obtain meaningful results, you must filter the following dimensions to child categories:

◆ Group
◆ Job Type
◆ Job Status

Otherwise, the data is averaged together for all the child categories of each dimension, which is not meaningful.

For example, suppose that in the Group dimension, you want to organize jobs statistics by project. You can do one of the following:

◆ Click the **Project** hypertext link in the row heading.

◆ In the Dimension line, click **Group > Project**.

Measures

The Daily Job Statistics cube includes the following measures:

Avg Job Number   The average number of jobs, based on a 15-minute default collection interval that is aggregated hourly. This is the average, based on weight by the sampling count measure (below). Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Min Job Number   The minimum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. This is the minimum, during hourly aggregation, of all the minimum values collected during each collection period. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Max Job Number   The maximum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. This is the maximum, during hourly aggregation, of all the maximum values collected during each collection period. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

How the Summary is reached

You may notice that the Total Avg number of jobs in the Summary line of the Daily Job Statistics report does not match the what is stated in the Summary line of the Daily Job Statistics cube. Both calculations are correct; however, the methods used to calculate an average are different.

Due to a different drill-down and roll-up for a cube's multiple dimensions, a weighted average is used (divide by weight average), while the report summary uses the normal average (divide by number of jobs).

Cognos Cube Views

The Daily Job Statistics cube does not include any Cognos cube views.

12

The Daily Job Throughput Cube

The Daily Job Throughput cube lets you explore the job throughput of different users and projects. The cube contains historical job throughput data. It is based on data collected from the `lsb.acct` file, collected in the Daily Job Throughput datamart. The dynamic data (from the Daily Job Throughput datamart) is collected hourly and aggregated daily. You can explore the job throughput for all queues, projects, users, and hosts.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Job Throughput cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The date when the job finished running, based on the date and time according to the LSF master host.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆   [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆   [Yesterday]
- ◆   WTD | [Day]
- ◆   MTD | [Week] | [Day]
- ◆   QTD | [Month] | [Week] | [Day]
- ◆   YTD | |[Month] | [Week] | [Day]
- ◆   [Previous Week] | [Day]
- ◆   [Previous Month] | [Week] | [Day]
- ◆   [Previous Quarter] | [Month] | [Week] | [Day]
- ◆   [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆   Sunday
- ◆   Monday
- ◆   Tuesday
- ◆   Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Queue	◆ [Priority] | [Level]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Project	◆ [User-defined project]
User	◆ [User] | [User]
# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4

# of Cores per CPU   The number of cores in each physical processor in the host.

- 1
- 2
- 4

# of Threads per Core   The maximum number of threads in each core in the host.

- 1
- 2

Measures

Data collection

Data about job throughput is collected in "real time", so that it does not include data about running or pending jobs. Therefore, jobs that were submitted before midnight but are still running or pending at the last midnight before the cube build are not included in the count of submitted jobs. Consequently, the data in the cubes data changes frequently, especially when there are lots of long jobs, which run for more than one day.

Consider the example of a single cluster that is monitored by Analytics with the following average hourly statistics:

- 300 jobs are submitted
- 200 jobs finish running with done status
- 50 jobs finish running with exit status
- 20 jobs never run
- 30 jobs are still running when Analytics analyzes the data

The average number of submitted jobs in the resulting cube will increase from the time Analytics originally analyzes the data to once the 30 running jobs finish running.

Calculation of measures

The measures in this cube are the total number of jobs during the specified period. You can drill into the data to see the hourly totals.

The Job Throughput cube includes the following measures:

of submitted jobs

Total number of submitted jobs.

of done jobs

Total number of done jobs.

of exited jobs

Total number of exited jobs.

Cognos Cube Views

The Daily Job Throughput cube includes the following Cognos cube view report:

- Overall job throughput by quarter

13

The Daily Jobs by Pending Reason Cube

The Daily Jobs by Pending Reason cube lets you explore the reasons jobs are pending. The cube contains historical data about the reasons jobs are pending. It is based on the Daily Jobs by Pending Reason datamart. You can explore these jobs by pending reason, type of pending reason, and cluster.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Other notes about the Group dimension

In order to obtain meaningful data, you must specify child categories (i.e. filter) for the Group dimension. You can filter this dimension from the Dimension line. For additional information, see "Filtering data" on page 79.

Dimensions

The Daily Jobs By Pending Reason cube includes the following dimensions and child categories:

Group   To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 79.

- ◆ By Pending Reason
  - ❖ Any pending reason in LSF
- ◆ By Pending Reason Type
  - ❖ Any pending reason type in LSF, such as User-Related Reasons, Job-Related Reasons, and Host-Related Reasons

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]

- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Filtering data

In order to obtain meaningful results, you must filter the Group dimension to a child category. In other words, you need to specify whether to display the data by project, queue, user, pending reason, or cluster. Otherwise, the data is averaged for all these groups together, which is not meaningful.

For example, suppose that in the Group dimension, you want to organize jobs statistics by project. You can do one of the following:

- ◆ Click the **By Project** hypertext link in the row heading.
- ◆ In the Dimension line, click **Group > By Project**.

Measures

The Daily Jobs By Pending Reason cube includes the following measures:

Avg Job Number
The average number of jobs, based on a 15-minute default collection interval that is aggregated hourly. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Min Job Number
The minimum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Max Job Number
The maximum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Cognos Cube Views

The Daily Jobs By Pending Reason cube does not include any Cognos cube views.

14

The Daily License Consumption Cube

The Daily License Consumption cube lets you explore up to six months of daily license consumption and utilization, to ensure that the license service level is good enough in your cluster. It contains information about license vendors, license servers, user groups, and execution host groups. It is based on the Daily License Consumption datamart.

Contents

Dimensions

Abbreviations

WTD Week to Date

MTD Month to Date

QTD Quarter to Date

YTD Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Consumption cube includes the following dimensions and child categories:

License vendor

- All vendors | [License vendor] | [License feature]

Date The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users| [User]
Host	◆ All Hosts | [Host]
License server	◆ [License Server]

Measures

The Daily License Consumption cube includes the following measures:

License minutes (number x minutes)

The total number of minutes that all licenses are used.

Cost

The dollar cost of a specific to a user, project, etc. Cost is calculated as **License minutes × License Rate**. You can enter the license rate in the console under **Clusters > [*Cluster Name*] > Resource Costs > License Rate**. You can specify the license rate as well as the start date for applying this rate.

Cognos Cube Views

The Daily License Consumption cube includes the following Cognos cube views:

- Consumption per server
- Consumption per weekday
- Users consuming the most licenses
- Weekly trend of consumption per feature

15

The Daily License Denial Cube

The Daily License Denial cube lets you explore the license denial and utilization, and to ensure that the license service level is good enough in your cluster. It contains information about license vendors, license servers, user groups, and execution host groups. It is based on the Daily License Denial datamart.

The data in the Daily License Denial cube is based on the FLEXlm debug log files. If you disable the creation of these debug log files in the Options file, the License Denial cube will be empty. For additional information about the FLEXlm debug log files, refer to the *FLEXlm End Users Guide*.

◆ This cube samples data frequently.

◆ Only recent data is available; data expires after 30 days.

◆ The data in this cube is aggregated hourly.

Differences between the License Denial cubes and the License Denial by Feature cubes

License Denial cubes	License Denial by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers
◆ cannot identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is identified	◆ can identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is recognized as false and ignored

Configuration In order for the license server name to be displayed in license denial cubes and Cognos cube views, you must configure the license log file in the following format:

```
Fullpath/license_log_file | License_server_name
```

Contents
◆ "Dimensions" on page 90
◆ "Measures" on page 92
◆ "Cognos Cube Views" on page 93

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Denial cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours.
User	◆ All Users | [User]
Host	◆ All Hosts | [Host]
License server	◆ [License Server]

Measures

The Daily License Denial cube includes the following measures:

License denials

The total number of license denials.

Cognos Cube Views

The Daily License Denial cube includes the following Cognos cube view report:

- Weekly trend of license details

16

The Daily License Denial By Feature Cube

The Daily License Denial By Feature cube lets you explore the historical data about global-wide license denial by feature and the service level of license servers. It contains information about license vendors, users, and execution hosts. It is based on the Daily License Denial By Feature datamart.

The data in the Daily License Denial By Feature cube is based on the FLEXlm debug log files. If you disable the creation of these debug log files in the Options file, the License Denial cube will be empty. For additional information about the FLEXlm debug log files, refer to the *FLEXlm End Users Guide*.

◆ This cube samples data frequently.

◆ Only recent data is available; data expires after 30 days.

◆ The data in this cube is aggregated hourly.

Differences between the License Denial cubes and the License Denial by Feature cubes

License Denial cubes	License Denial by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers
◆ cannot identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is identified	◆ can identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is recognized as false and ignored

Configuration   In order for the license server name to be displayed in license denial cubes and Cognos cube views, you must configure the license log file in the following format:

```
Fullpath/license_log_file | License_server_name
```

Contents   ◆ "Dimensions" on page 96

◆ "Measures" on page 98

◆ "Cognos Cube Views" on page 99

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Denial by Feature cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The date when the job finished running, based on the date and time according to the LSF master host.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday

◆ Saturday

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours
◆ Non-core hours

User ◆ [User] | [User]

Host ◆ All Hosts | [Host]

Measures

The Daily License Denial by Feature cube includes the following measures:

License denials

The total number of license denials.

Cognos Cube Views

The Daily License Denial By Feature cube does not include any Cognos cube views.

C H A P T E R

17

The Daily License Usage By User Cube

The Daily License Usage By User cube lets you explore up to six months of data about daily license usage from a user perspective. It contains information about license vendors, license servers, user groups, and execution host groups. This cube is based on the Daily License Usage datamart. For information about license utilization from a license server perspective, see "The Daily License Utilization Cube" on page 125.

Contents
- ◆ "Dimensions" on page 102
- ◆ "Measures" on page 104
- ◆ "Cognos Cube Views" on page 105

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Usage By User cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

User
- All Users | [User]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday

	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours.
License server	◆ All License Servers | [License Server]

Measures

The Daily License Usage By User cube includes the following measures:

Average usage by user

The average license usage by user. Usage is expressed as a number.

Peak usage by user

The maximum license usage by user. Usage is expressed as a number.

Average # of available licenses

The average total number of available licenses in a specified time period.

Cognos Cube Views

The Daily License Usage By User cube includes the following Cognos cube views:

- Peak and average usage by user
- User share by feature

18

The Daily License Usage By User Group Cube

The Daily License Usage By User Group cube lets you explore six months of data about daily license usage from a user group perspective. It contains information about license vendors, license servers, user groups, and execution host groups. This cube is based on the Daily License Usage By User Group datamart. For information about license utilization from a license server perspective, see "The Daily License Utilization Cube" on page 125.

Note: By default, Analytics does not build this cube. For information about building cubes, refer to *Administering Platform Analytics*.

Contents
- "Dimensions" on page 108
- "Measures" on page 110
- "Cognos Cube Views" on page 111

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Usage by User Group cube includes the following dimensions and child categories:

License vendor
: ◆ All vendors | [License vendor] | [License feature]

Usergroup
: ◆ All Usergroups| [Usergroup]

Date
: The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆ [Yesterday]
◆ WTD | [Day]
◆ MTD | [Week] | [Day]
◆ QTD | [Month] | [Week] | [Day]
◆ YTD | |[Month] | [Week] | [Day]
◆ [Previous Week] | [Day]
◆ [Previous Month] | [Week] | [Day]
◆ [Previous Quarter] | [Month] | [Week] | [Day]
◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday
: Based on the date and time according to the LSF master host.

◆ Sunday
◆ Monday
◆ Tuesday
◆ Wednesday
◆ Thursday

	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
License server	◆ All License Servers | [License Server]

Measures

The Daily License Usage By User Group cube includes the following measures:

Average usage by usergroup

The average license usage by user group. Usage is expressed as a number.

Peak usage by usergroup

The maximum license usage by user group. Usage is expressed as a number.

Average # of available licenses

The average total number of available licenses in a specified time period.

Cognos Cube Views

The Daily License Usage By User Group cube does not include any Cognos cube views.

19

The Daily License Usage Statistics by Feature Cube

The Daily License Usage Statistics by Feature cube lets you explore up to six months of daily license usage statistics. It contains information about license denial, utililization, and consumption, all by feature in one cube. It is based on the Daily License Usage Statistics by Feature - Denial, Daily License Usage Statistics by Feature - Utilization, and Daily License Usage Statistics by Feature - Consumption datamarts.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Usage Statistics by Feature cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users| [User]
Host	◆ All Hosts | [Host]

Measures

The Daily License Usage Statistics by Feature cube includes the following measures:

License denials

The total number of license denials.

Average usage by feature

The average license usage by feature. Usage is expressed as a number.

Peak usage by feature

The maximum license usage by feature. Usage is expressed as a number.

Average utilization by feature(%)

The average license utilization for each feature over a period of time, expressed as a percentage.

Average # of available licenses

The average total number of available licenses in a specified time period.

Peak # of available licenses by feature

The peak number of licenses available for each feature over a period of time.

Peak utilization by feature (%)

The peak license utilization for each feature over a period of time, expressed as a percentage.

# of reserved licenses by feature

The number of reserved licenses for each feature.

Utilization at peak usage by feature (%)

The peak license utilization for each feature at the time when the maximum number of licenses are being used.

License available minutes

The amount of time that licenses were available, in minutes. This is calculated as **number of licenses_available** $\times$ **minutes available**.

License minutes

The amount of time that licenses were used, in minutes. This is calculated as **number of licenses_used** $\times$ **minutes used**.

Cost

Cost is calculated as **minutes_used** $\times$ **Resource Rate**. The Resource Rate is defined in the console under **Clusters > [*Cluster Name*] > Resource Costs > License Cost Rate**. You can specify the job cost rate as well as the start date for applying this rate.

Cognos Cube Views

The Daily License Usage Statistics by Feature cube does not include any Cognos cube views.

20

The Daily License Usage Statistics Cube

The Daily License Usage Statistics cube lets you explore up to six months of daily license usage statistics. It contains information about license denial, utililization, and consumption in one cube. It is based on the Daily License Usage Statistics - Denial, Daily License Usage Statistics - Utilization, and Daily License Usage Statistics - Consumption datamarts.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Usage Statistics cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date
: The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday
: Based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users| [User]
Host	◆ All Hosts	[Host]
License server	◆ All License Servers	[License Server]

Measures

The Daily License Usage Statistics cube includes the following measures:

License denials

The total number of license denials.

Average usage by server

The average license usage by server. Usage is expressed as a number.

Peak usage by server

The maximum license usage by server. Usage is expressed as a number.

Average utilization by server (%)

The average license utilization on each license server over a period of time, expressed as a percentage.

Peak utilization by server (%)

The peak license utilization on each license server over a period of time, expressed as a percentage.

Average # of available licenses

The average total number of available licenses in a specified time period.

Peak # of available licenses by server

The peak number of licenses available on each license server over a period of time.

Utilization at peak usage by server (%)

The peak license utilization on each license server at the time when the maximum number of licenses are being used.

# of reserved licenses by server

The number of reserved licenses on each license server.

License available minutes

The amount of time that licenses were available, in minutes. This is calculated as **number of licenses_available × minutes available**.

License minutes

The amount of time that licenses were used, in minutes. This is calculated as **number of licenses_used × minutes used**.

Cost

Cost is calculated as **minutes_used × Resource Rate**. The Resource Rate is defined in the console under **Clusters > [*Cluster Name*] > Resource Costs > License Cost Rate**. You can specify the job cost rate as well as the start date for applying this rate.

Cognos Cube Views

The Daily License Usage Statistics cube does not include any Cognos cube views.

21

The Daily License Utilization Cube

The Daily License Utilization cube lets you explore license utilization from the license server perspective. It contains up to six months of historical license utilization data, and it is based on the Daily License Utilization datamart. It does not contain information about individual users. For information about license usage from a user perspective, see "The Daily License Usage By User Cube" on page 101.

Differences between the License Utilization cubes and the License Utilization by Feature cubes

License Utilization cubes	License Utilization by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Utilization cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

License server
- [License Server]

Measures

The Daily License Utilization cube includes the following measures:

Average usage by server

The average number of licenses in use on each license server.

Peak usage by server

The maximum number of licenses in use on each license server.

Average utilization by server (%)

The average license utilization on each license server over a period of time, expressed as a percentage.

Peak utilization by server (%)

The peak license utilization on each license server over a period of time, expressed as a percentage.

Average # of available licenses

The total number of licenses available.

Peak # of available licenses by server

The peak number of licenses available on each license server over a period of time.

Utilization at peak usage by server (%)

The peak license utilization on each license server at the time when the maximum number of licenses are being used.

of reserved licenses by server

The number of reserved licenses on each license server.

Cognos Cube Views

The Daily License Utilization cube includes the following Cognos cube views:

- \# of licenses available on each server
- Features ranked by peak utilization
- License usage and available licenses by server
- Overall utilization by feature
- Overall utilization by vendor and server
- Weekly trend of utilization per feature

22

The Daily License Utilization By Feature Cube

The Daily License Utilization By Feature cube lets you explore license utilization from the license server perspective. It contains up to six months of historical license utilization data by feature, and it is based on the Daily License Utilization by Feature datamart. It does not contain information about individual users. For information about license usage from a user perspective, see "The Daily License Usage By User Cube" on page 101.

Differences between the License Utilization cubes and the License Utilization by Feature cubes

License Utilization cubes	License Utilization by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily License Utilization By Feature cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Measures

The Daily License Utilization By Feature cube includes the following measures:

Average usage by feature

The average number of licenses in use for each feature.

Peak usage by feature

The maximum number of licenses in use for each feature.

Average utilization by feature (%)

The average license utilization for each feature over a period of time, expressed as a percentage.

Average # of available licenses

The average number of licenses available on the license server.

Peak # of available licenses by feature

The peak total number of licenses available for each feature over a period of time.

Peak utilization by feature (%)

The peak license utilization for each feature over a period of time, expressed as a percentage.

Utilization at peak usage by feature (%)

The utilization for each feature during peak usage, expressed as a percentage.

of reserved licenses by feature

The number of licenses reserved for each feature.

Cognos Cube Views

The Daily License Utilization By Feature cube does not include any Cognos cube views.

23

The Daily Shared Resource Cube

The Daily Shared Resource cube contains historical data about the percentage of utilization and usage of shared resources. It shows only numeric, dynamic, user-defined resources in LSF. Information about shared resources is based on data collected from LSF, in the Daily Shared Resource Usage datamart.

◆ The data in this cube is aggregated hourly.

◆ The data in this cube expires after 30 days.

Contents ◆ "Dimensions" on page 134
◆ "Measures" on page 136
◆ "Cognos Cube Views" on page 137

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Shared resource
A cluster-based LSF resource that is shared among hosts, e.g. a license.

Cluster
A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Shared Resource Usage cube includes the following dimensions and child categories:

Resource
[All Shared Resources] | [Shared Resource]

Cluster
Cluster | [Host]

Date
The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Business Hour
Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours

- ◆ Non-core hours

Measures

In order to obtain meaningful data for the **Average usage**, **Average utilization (%)**, and **Average total available** measures (described below), you must make sure that there is data in the LIC_FEATURE_NAME and TOTAL_AVAIL columns in the WI_LICENSETOTAL table. In general, this table contains data from the LICENSEUSAGE table. However, if the you disabled this table, or if the shared resources are not licenses, you need enter the following information in the WI_LICENSETOTAL table:

◆ shared resource name in the LIC_FEATURE_NAME column
◆ total number of shared resources in the TOTAL_AVAIL column

You can enter this information in one of the following ways:

◆ use a database tool
◆ use a script, such as SQL

Average usage

The average number of shared resources.

Average utilization (%)

The average utilization of shared resources by server, expressed as a percentage.

Average total available

The average total number of available shared resource, including licenses in use and licenses that are free for use.

Average free

The average total number of licenses that are free for use.

Cognos Cube Views

The Daily Shared Resource cube does not include any Cognos cube views.

24

The Daily Workload Cube

The Daily Workload cube lets you explore how much hardware and time different users and projects consume, and discover whether cluster usage actually maps to your business goals and strategies. It contains historical workload data about jobs, queues, projects, and hosts. It is based on data collected from the `lsb.acct` file, collected in the Daily Job Information datamart. You can explore the workload for all projects, queues, users, and hosts.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Workload cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The date when the job finished running, based on the date and time according to the LSF master host.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   The weekday when the job was finished.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Date (Submit)	The date when the job was submitted, based on the date and time according to the LSF master host.

◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆ QTD | [Month] | [Week] | [Day]
◆ Prior QTD | [Month] | [Week] | [Day]
◆ QTD Grouped
  ❖ Prior QTD | [Month] | [Week] | [Day]
  ❖ QTD | [Month] | [Week] | [Day]
  ❖ QTD Change
  ❖ QTD Growth
◆ YTD | [Quarter] | [Month] | [Week] | [Day]
◆ Prior YTD | [Quarter] | [Month] | [Week] | [Day]
◆ YTD Grouped
  ❖ Prior YTD | [Quarter] | [Month] | [Week] | [Day]
  ❖ YTD | [Quarter] | [Month] | [Week] | [Day]
  ❖ YTD Change
  ❖ YTD Growth
◆ MTD | [Week] | [Day]
◆ Prior MTD | [Week] | [Day]
◆ MTD Grouped
  ❖ Prior MTD | [Week] | [Day]
  ❖ MTD | [Week] | [Day]
  ❖ MTD Change
  ❖ MTD Growth
◆ WTD | [Day]
◆ Prior WTD | [Day]
◆ WTD Grouped
  ❖ Prior WTD | [Day]
  ❖ WTD | [Day]
  ❖ WTD Change
  ❖ WTD Growth

Weekday (Submit) The weekday when the job was submitted, based on the date and time according to the LSF master host.

◆ Sunday
◆ Monday

- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Date/Start   The date when the job started running, based on the date and time according to the LSF master host.

- No start date
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]
- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- Prior QTD | [Month] | [Week] | [Day]
- QTD Grouped
    - Prior QTD | [Month] | [Week] | [Day]
    - QTD | [Month] | [Week] | [Day]
    - QTD Change
    - QTD Growth
- YTD | [Quarter] | [Month] | [Week] | [Day]
- Prior YTD | [Quarter] | [Month] | [Week] | [Day]
- YTD Grouped
    - Prior YTD | [Quarter] | [Month] | [Week] | [Day]
    - YTD | [Quarter] | [Month] | [Week] | [Day]
    - YTD Change
    - YTD Growth
- MTD | [Week] | [Day]
- Prior MTD | [Week] | [Day]
- MTD Grouped
    - Prior MTD | [Week] | [Day]
    - MTD | [Week] | [Day]
    - MTD Change
    - MTD Growth
- WTD | [Day]
- Prior WTD | [Day]
- WTD Grouped
    - Prior WTD | [Day]
    - WTD | [Day]
    - WTD Change
    - WTD Growth

Weekday/Start   The weekday when the job started running, based on the date and time according to the LSF master host.

- No start date
- Sunday

- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Project
- [User-defined project]

User
- [All Users] | [User]

Queue
- [Priority] | [Level]

Application Tag
- [Application Tag]

Host Type   Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
- [Execution Host Type] | [Host]

Host Model   Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
- [Execution Host Model] | [Host]

# of CPUs per Host   The number of physical processors in the execution host.

- 1
- 2
- 4

# of Cores per CPU   The number of cores in each physical processor in the host.

- 1
- 2
- 4

# of Threads per Core   The maximum number of threads in each core in the host.

- 1
- 2

Hosts (Submit)   Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Submission Host Type] | [Host]

Job Type
- [User-defined job]

# of CPUs per Job   The number of CPUs used to run the job. This is not necessarily the same as the number of CPUs on the host.

- 0.0 (The job was never executed.)
- 1.0
- 4.0

Job Exit Status
- DONE
- EXIT

Job Exit Code
- [Job Exit Code]

Pending Time Rank
- 0 sec to < 5 sec
- 5 sec to < 30 sec
- 30 sec to < 5 min
- 5 min to < 60 min
- 60 min to < 1 day
- One day or more

Run Time Rank
- 0 sec to < 1 min
- 1 min to < 5 min
- 5 min to < 20 min
- 20 min to < 60 min
- 60 min to < 4 hour
- 4 hour to < 1 day
- One day or more

CPU Time Rank
- 0 sec to < 1 min
- 1 min to < 5 min
- 5 min to < 20 min
- 20 min to < 60 min

◆ 60 min to < 4 hour

◆ 4 hour to < 1 day

◆ One day or more

Memory Usage Rank ◆ 0 GB to < 1 GB

◆ 1 GB to < 2 GB

◆ 2 GB to < 4 GB

◆ 4 GB to < 8 GB

◆ 8 GB to < 16 GB

◆ 16 GB to < 32 GB

◆ 32 GB or more

User Group Name of user group as defined in Platform LSF.

◆ [User Group]

Note: If a user belongs to multiple user groups, the job submitted by that user is not associated with a group unless "-G"' is used to indicate the group when submitting the job.

Job ID The Job ID as assigned in Platform LSF.

◆ [Job ID]

Note: To use this dimension, you must enable the "Include Job ID in Workload Cubes" feature in the Platform Analytics Console (click **Addons** in the navigation tree).

Measures

The Daily Workload cube includes the following measures:

# of Jobs

Total number of jobs.

Job run time (min)

Total run time of all jobs, in minutes.

Job pend time (min)

Total pend time of all jobs, in minutes.

Job turnaround time (min)

Total time elapsed between job submission and job finish.

CPU time (sec)

Total CPU time used by all jobs, in seconds.

Used mem (MB)

Total memory used by all jobs, in megabytes.

Used swap (MB)

Total swap space used by all jobs, in megabytes.

CPU time/Job run time (%)

Total real CPU time used by a job, divided by the total time elapsed between job start and job finish.

Avg run time/job (min)

Average run time of a job, in minutes.

Avg pend time/job (min)

Average pend time of a job, in minutes.

Avg turnaround time/job (min)

Average time elapsed between job submission and job finish for all specified jobs.

Job run time/Turnaround time (%)

Total time elapsed between job start and job finish, divided by the total time elapsed between job submission and job finish.

Avg CPU time/job (min)

Average CPU time used by a job, in minutes.

Avg used mem/job (MB)

Average memory used by a job, in megabytes.

Avg used swap/job (MB)

Average swap space used by a job, in megabytes.

Cognos Cube Views

The Daily Workload cube includes the following Cognos cube view report:

- Quarterly workload by cluster

25

The Daily Workload License Cube

The Daily Workload License cube lets you explore the daily historical workload license data obtained from your LSF cluster and the FLEXnet Manager server. This cube is based on the Daily Workload License datamart.

Note: Macrovision's FLEXnet Manager is a central management tool for all FLEXlm servers in an enterprise environment. In order to build and explore the data that makes up this cube, you require a FLEXnet Manager installation.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Daily Workload License cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Date (Submit)   The date when the job was submitted of which you want sampled, based on the date and time according to the LSF master host.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ Current Day
- ◆ Yesterday
- ◆ WTD | [Day]
- ◆ Prior WTD | [Day]
- ◆ WTD Grouped
  - ❖ Prior WTD | [Day]
  - ❖ WTD | [Day]
  - ❖ WTD Change
  - ❖ WTD Growth
- ◆ MTD | [Week] | [Day]
- ◆ Prior MTD | [Week] | [Day]
- ◆ MTD Grouped
  - ❖ Prior MTD | [Week] | [Day]
  - ❖ MTD | [Week] | [Day]
  - ❖ MTD Change
  - ❖ MTD Growth
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ Prior QTD | |[Month] | [Week] | [Day]
- ◆ QTD Grouped
  - ❖ Prior QTD | [Month] | [Week] | [Day]
  - ❖ QTD | [Month] | [Week] | [Day]
  - ❖ QTD Change
  - ❖ QTD Growth
- ◆ YTD | [Quarter] | [Week] | [Day]
- ◆ YTD Grouped
  - ❖ Prior YTD | [Quarter] | [Week] | [Day**]**
  - ❖ YTD | [Quarter] | [Week] | [Day]
  - ❖ YTD Change
  - ❖ YTD Growth

Weekday (Submit)   The weekday when the job was submitted of which you want sampled, based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Date/Start   The date when the job started running of which you want sampled, based on the date and time according to the LSF master host.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ Current Day
- ◆ Yesterday
- ◆ WTD | [Day]
- ◆ Prior WTD | [Day]
- ◆ WTD Grouped
    - ❖ Prior WTD | [Day]
    - ❖ WTD | [Day]
    - ❖ WTD Change
    - ❖ WTD Growth
- ◆ MTD | [Week] | [Day]
- ◆ Prior MTD | [Week] | [Day]
- ◆ MTD Grouped
    - ❖ Prior MTD | [Week] | [Day]
    - ❖ MTD | [Week] | [Day]
    - ❖ MTD Change
    - ❖ MTD Growth
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ Prior QTD | |[Month] | [Week] | [Day]
- ◆ QTD Grouped
    - ❖ Prior QTD | [Month] | [Week] | [Day]
    - ❖ QTD | [Month] | [Week] | [Day]
    - ❖ QTD Change
    - ❖ QTD Growth
- ◆ YTD | [Quarter] | [Week] | [Day]
- ◆ YTD Grouped
    - ❖ Prior YTD | [Quarter] | [Week] | [Day]
    - ❖ YTD | [Quarter] | [Week] | [Day]
    - ❖ YTD Change
    - ❖ YTD Growth

Weekday/Start   The weekday when the job started running of which you want sampled, based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday

	◆ Friday
	◆ Saturday
License Server	◆ [License Server]
Project	◆ [User-defined Project]
User	◆ All Users	[User]
Queue	◆ [Priority]	[Level]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type]	[Host]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model]	[Host]
Hosts (Submit)	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Submission Host Type]	[Host]
Job Group	Name of the job group as defined in Platform LSF.
	◆ [Job Group]
# of CPUs	The number of CPUs used to run the job. This is not necessarily the same as the number of CPUs on the host.
	◆ 0.0 (The job was never executed.)
	◆ 1.0
	◆ 4.0
Job Exit Status	◆ DONE
	◆ EXIT
Job Exit Code	◆ [Job Exit Code]
Pending Time Rank	◆ 0 sec to < 5 sec
	◆ 5 sec to < 30 sec
	◆ 30 sec to < 5 min
	◆ 5 min to < 60 min
	◆ 60 min to < 1 day
	◆ One day or more
Vendor	◆ All vendors	[License vendor]	[License feature]
Event	◆ DENIED
	◆ USED
Event Reason	◆ Licensed number of users already reached
	◆ Normal
Job ID	The Job ID as assigned in Platform LSF.
	◆ [Job ID]

Note: To use this dimension, you must enable the "Include Job ID in Workload Cubes" feature in the Platform Analytics Console (click **Addons** in the navigation tree).

Measures

The Daily Workload License cube includes the following measures:

# of jobs

Total number of jobs.

Job run time (min)

Total run time of all jobs, in minutes.

Job pend time (min)

Total pend time of all jobs, in minutes.

Job turnaround time (min)

Total time elapsed between job submission and job finish.

CPU time (sec)

Total CPU time used by all jobs, in seconds.

Used mem (MB)

Total memory used by all jobs, in megabytes.

Used swap (MB)

Total swap space used by all jobs, in megabytes.

CPU time/Job run time (%)

Total real CPU time used by a job, divided by the total time elapsed between job start and job finish.

Avg run time/job (min)

Average run time of a job, in minutes.

Avg pend time/job (min)

Average pend time of a job, in minutes.

Avg turnaround time/job (min)

Average time elapsed between job submission and job finish for all specified jobs.

Job run time/Turnaround time (%)

Total time elapsed between job start and job finish, divided by the total time elapsed between job submission and job finish.

Avg CPU time/job (min)

Average CPU time used by a job, in minutes.

Avg used mem/job (MB)

Average memory used by a job, in megabytes.

Avg used swap/job (MB)

Average swap space used by a job, in megabytes.

# of checkouts

Number of license checkouts.

Max # of checkouts

Top number of license checkouts at one time for each job; measures the frequency of checkouts for a single application.

Total feature time (min)

Total license feature time per job, in minutes; combines all consumed license time.

Avg feature time (min)

Average license feature time per job, in minutes; averages the consumed license time.

of events

Number of events (for example, denial, support, check in, check out).

Cognos Cube Views

The Daily Workload License cube does not include any Cognos cube views.

26

The Daily HPC Workload Cube

The HPC Daily Workload cube lets you explore how much hardware and time different users and projects consume, and discover whether cluster usage actually maps to your business goals and strategies. It contains historical workload data about HPC jobs, queues, projects, and hosts. It is based on data collected in the Host Properties and HPC Job Information datamarts. You can explore the workload for all HPC projects, queues, users, and hosts.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Daily Workload cube includes the following dimensions and child categories:

Job
- [Cluster] | [Name] | [Job Type]

Date   The date when the job finished running, based on the date and time according to the LSF master host.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   The weekday when the job was finished.

- Sunday
- Monday
- Tuesday
- Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Date (Submit)	The date when the job was submitted, based on the date and time according to the LSF master host.
	◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
	◆ QTD | [Month] | [Week] | [Day]
	◆ Prior QTD | [Month] | [Week] | [Day]
	◆ QTD Grouped
	❖ Prior QTD | [Month] | [Week] | [Day]
	❖ QTD | [Month] | [Week] | [Day]
	❖ QTD Change
	❖ QTD Growth
	◆ YTD | [Quarter] | [Month] | [Week] | [Day]
	◆ YTD Grouped
	❖ Prior YTD | [Quarter] | [Month] | [Week] | [Day]
	❖ YTD | [Quarter] | [Month] | [Week] | [Day]
	❖ YTD Change
	❖ YTD Growth
	◆ MTD | [Week] | [Day]
	◆ Prior MTD | [Week] | [Day]
	◆ MTD Grouped
	❖ Prior MTD | [Week] | [Day]
	❖ MTD | [Week] | [Day]
	❖ MTD Change
	❖ MTD Growth
	◆ WTD | [Day]
	◆ Prior WTD | [Day]
	◆ WTD Grouped
	❖ Prior WTD | [Day]
	❖ WTD | [Day]
	❖ WTD Change
	❖ WTD Growth
Weekday (Submit)	The weekday when the job was submitted, based on the date and time according to the LSF master host.
	◆ Sunday
	◆ Monday
	◆ Tuesday

- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Date/Start   The date when the job started running, based on the date and time according to the LSF master host.

- ◆ No start date
- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ Prior QTD | [Month] | [Week] | [Day]
- ◆ QTD Grouped
  - ❖ Prior QTD | [Month] | [Week] | [Day]
  - ❖ QTD | [Month] | [Week] | [Day]
  - ❖ QTD Change
  - ❖ QTD Growth
- ◆ YTD | [Quarter] | [Month] | [Week] | [Day]
- ◆ YTD Grouped
  - ❖ Prior YTD | [Quarter] | [Month] | [Week] | [Day]
  - ❖ YTD | [Quarter] | [Month] | [Week] | [Day]
  - ❖ YTD Change
  - ❖ YTD Growth
- ◆ MTD | [Week] | [Day]
- ◆ Prior MTD | [Week] | [Day]
- ◆ MTD Grouped
  - ❖ Prior MTD | [Week] | [Day]
  - ❖ MTD | [Week] | [Day]
  - ❖ MTD Change
  - ❖ MTD Growth
- ◆ WTD | [Day]
- ◆ Prior WTD | [Day]
- ◆ WTD Grouped
  - ❖ Prior WTD | [Day]
  - ❖ WTD | [Day]
  - ❖ WTD Change
  - ❖ WTD Growth

Weekday/Start   The weekday when the job started running, based on the date and time according to the LSF master host.

- ◆ No start date
- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Project	◆ [User-defined project]
User	◆ [All Users] | [User]
Queue	◆ [Priority] | [Level]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Hosts (Submit)	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Submission Host Type] | [Host]
Job Type	◆ [User-defined job]
# of CPUs	The number of CPUs used to run the job. This is not necessarily the same as the number of CPUs on the host.
	◆ 0.0 (The job was never executed.)
	◆ 1.0
	◆ 4.0
Job Exit Status	◆ [Job Type] | [Job Exit Status]
Job Exit Code	◆ [Job Type] | [Job Exit Code]
Pending Time Rank	◆ 0 sec to < 5 sec
	◆ 5 sec to < 30 sec
	◆ 30 sec to < 5 min
	◆ 5 min to < 60 min
	◆ 60 min to < 1 day
	◆ One day or more
User Group	Name of user group as defined in Platform LSF.
	◆ [User Group]

Note: If a user belongs to multiple user groups, the job submitted by that user is not associated with a group unless "-G" is used to indicate the group when submitting the job.

# of CPUs per Host	The number of physical processors in the execution host.
	◆ 1
	◆ 2
	◆ 4

# of Cores per CPU   The number of cores in each physical processor in the host.

  ◆   1
  ◆   2
  ◆   4

# of Threads per Core   The maximum number of threads in each core in the host.

  ◆   1
  ◆   2

Application Tag   ◆   [Application Tag]

Measures

The Daily Workload cube includes the following measures:

Job run time (min)

Total run time of all jobs, in minutes.

Job pend time (min)

Total pend time of all jobs, in minutes.

Job turnaround time (min)

Total time elapsed between job submission and job finish.

CPU time (sec)

Total CPU time used by all jobs, in seconds.

Used mem (MB)

Total memory used by all jobs, in megabytes.

Used swap (MB)

Total swap space used by all jobs, in megabytes.

CPU time/Job run time (%)

Total real CPU time used by a job, divided by the total time elapsed between job start and job finish.

Avg run time/job (min)

Average run time of a job, in minutes.

Avg pend time/job (min)

Average pend time of a job, in minutes.

Avg turnaround time/job (min)

Average time elapsed between job submission and job finish for all specified jobs.

Job run time/Turnaround time (%)

Total time elapsed between job start and job finish, divided by the total time elapsed between job submission and job finish.

Avg CPU time/job (min)

Average CPU time used by a job, in minutes.

Avg used mem/job (MB)

Average memory used by a job, in megabytes.

Avg used swap/job (MB)

Average swap space used by a job, in megabytes.

Cognos Cube Views

The Daily Workload cube includes the following Cognos cube view report:

◆ Quarterly workload by cluster

27

The Hardware Performance Prediction Cube

The Hardware Performance Prediction cube contains historical and predictive data about hardware performance. It is based on the Host Properties and Hardware Performance Prediction datamarts.

Contents

Dimensions

Definitions

Change
The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster
A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hardware Performance Prediction cube includes the following dimensions and child categories:

Metric
Any metric in this cube. For more information, see "Metrics" on page 170.

Weekday
The predicted weekday.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Cluster
- Cluster | [Host]

Host Type
Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Type] | [Host]

Host Model
Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Model] | [Host]

Activity Level
- Critical
- Normal
- Under-utilized
- Not specified

Metric value categorizations for today.

Activity Level After 4 Weeks
- Critical
- Normal
- Under-utilized
- Not specified

Metric value categorizations for 4 weeks from now.

Activity Level After 8 Weeks

- Critical
- Normal
- Under-utilized
- Not specified

Metric value categorizations for 8 weeks from now.

Activity Level After 12 Weeks

- Critical
- Normal
- Under-utilized
- Not specified

Metric value categorizations for 12 weeks from now.

Measures Any measure in this cube. For more information, see "Measures" on page 172.

Metrics

The Hardware Performance Prediction cube includes the following metrics:

of logins

Description   Number of logins.

Internal name   ls

15min run queue

Description   Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Internal name   r15m

15s run queue

Description   Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

Internal name   r15s

1min run queue

Description   Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

Internal name   r1m

CPU (%)

Description   Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

Internal name   ut

Disk IO (KB/second)

Description   Average disk I/O, for all disks on a host, in kilobytes per second.

Internal name   io

Free mem (MB)

Description   Available physical memory, in megabytes.

Internal name   mem

Free swap (MB)

Description   Available swap space, in megabytes.

Internal name   swp

Free tmp (MB)

Description   Available space in a temporary file system, in megabytes.

| Internal name | tmp |

Inter. idle time (min)

| Description | Number of minutes since any interactive terminal activity. |
| Internal name | it |

Max mem (MB)

| Description | Average total amount of maximum physical memory, in megabytes. |
| Internal name | maxmem |

Max swap (MB)

| Description | Maximum available swap space, in megabytes. |
| Internal name | maxswp |

Max tmp (MB)

| Description | Maximum available space in a temporary file system, in megabytes. |
| Internal name | maxtmp |

Paging (pgs io/s)

| Description | Paging activity, in pages in and out per second. |
| Internal name | pg |

usedSlots/numCpus

| Description | The ratio between the used slots and the number of CPUs |

Measures

The Hardware Performance Prediction cube includes the following measures:

Average hourly value

The average of all average values in the history for this host and metric.

Average hourly maximum

The average of all maximum values in the history for this host and metric.

Average hourly minimum

The average of all minimum values in the history for this host and metric.

Maximum hourly value

The maximum of all maximum values in the historical data for this host and metric.

Minimum hourly value

The minimum of all minimum values in the historical data for this host and metric.

Number of systems

The number of hosts.

Metric critical threshold

The critical threshold for a specific metric. You can define the critical thresholds for metrics in the Metric Thresholds tab of the Options dialog box in the console. For additional information, see *Administering Platform Analytics*.

Times critical threshold exceeded

A count of the number of times in the historical data that the hourly maximum or minimum value of the metric value exceeded the critical threshold.

- For some metrics, such as **CPU (%)** (the percentage of CPU consumption), the greater the value is, the more critical the situation. Therefore, it is the hourly *maximum* value of this metric that could exceed the critical threshold.
- For other metrics, such as **Free Mem (MB)** (the amount of free real memory), the smaller the value is, the more critical the situation. Therefore, it is the hourly *minimum* value of this metric that could exceed the critical threshold.

You can view or modify metric thresholds in the Metric Thresholds tab of the Options dialog box.

By default, 24 weeks is the default date range of historical data that are checked for this calculation. You can modify the date range of checked historical value in the Extractor (JDBC type only) in the ETL tab of the Datamart Properties dialog box.

- NA is displayed for metrics without thresholds or for metrics whose critical thresholds were never exceeded in the historical data.

Days to critical threshold

The number of days before a metric reaches its critical threshold.

- NA is displayed for metrics without thresholds or for metrics whose critical thresholds will never be exceeded.
- 0 is displayed if the critical threshold was crossed in the past, or was predicted to have been crossed before the current day, including any gap between the end of the historical data and before the current day.

4th week forecasted average

Average value predicted for this host and metric in the 4th week.

4th week forecasted maximum

Maximum value predicted for this host and metric in 4 weeks.

4th week forecasted minimum

Minimum value predicted for this host and metric in 4 weeks.

8th week forecasted average

Average value predicted for this host and metric in the 8th week.

8th week forecasted maximum

Maximum value predicted for this host and metric in 8 weeks.

8th week forecasted minimum

Minimum value predicted for this host and metric in 8 weeks.

12th week forecasted average

Average value predicted for this host and metric in the 12th week.

12th week forecasted maximum

Maximum value predicted for this host and metric in 12 weeks.

12th week forecasted minimum

Minimum value predicted for this host and metric in 12 weeks.

Cognos Cube Views

The Hardware Performance Prediction cube includes the following Cognos cube view report:

- Performance trend over the next 12 weeks

28

The Hourly FLEX License Cube

The Hourly FLEX License cube lets you explore historical daily data about FLEX license usage and other events contained in the FLEXnet Manager server. This cube is based on the Hourly FLEX License datamart.

Note: Macrovision's FLEXnet Manager is a central management tool for all FLEXlm servers in an enterprise environment. In order to build and explore the data that makes up this cube, you require a FLEXnet Manager installation.

Contents
- "Dimensions" on page 176
- "Measures" on page 178
- "Cognos Cube Views" on page 179

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly FLEX License cube includes the following dimensions and child categories:

Cluster   Cluster | [Name]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hours	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.	
	◆	Core hours
	◆	Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.	
	◆	Shift 1 | [Hour]
	◆	Shift 2 | [Hour]
	◆	Shift 3 | [Hour]
	◆	By Hour | [Hour]
Project	◆	[User-defined Project]
License Server	◆	[License Server]
Vendor	◆	All vendors | [License vendor] | [License feature]
Version	◆	All versions | [License version][
LSF/non-LSF	◆	LSF
	◆	Non-LSF
Event	◆	DENIED
	◆	USED
Event Reason	◆	Licensed number of users already reached
	◆	Normal
	◆	UNKN
User	◆	[All Users] | [User]
Host	◆	All Hosts | [Host]

Measures

The Hourly FLEX License cube includes the following measures:

of checkouts

Number of license checkouts.

Max # of checkouts

Top number of license checkouts at one time for each job; measures the frequency of checkouts for a single application.

Total feature time (min)

Total license feature time per job, in minutes; combines all consumed license time.

Avg feature time (min)

Average license feature time per job, in minutes; averages the consumed license time.

of events

Number of events (for example, denial, used).

Cognos Cube Views

The Hourly FLEX License cube does not include any Cognos cube views.

29

The Hourly Hardware Cube

The Hourly Hardware cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Host Properties and Hourly Hardware datamarts. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, and host types by hour.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday

	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4

# of Cores per CPU   The number of cores in each physical processor in the host.

- 1
- 2
- 4

# of Threads per Core   The maximum number of threads in each core in the host.

- 1
- 2

Measures

The Hourly Hardware cube includes the following measures:

# of logins	Number of logins.
CPU (%)	Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.
Used mem (%)	Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).
Used swap (%)	Used swap space as a percentage.
Used tmp (%)	Used temp space as a percentage.
15s run queue	Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.
1min run queue	Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.
15min run queue	Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.
Paging (pgs io/s)	Paging activity, in pages in and out per second.
Disk IO (KB/second)	Average disk I/O, for all disks on a host, in kilobytes per second.
Inter. idle time (min)	Number of minutes since any interactive terminal activity.
Free mem (MB)	Available physical memory, in megabytes.
Free swap (MB)	Available swap space, in megabytes.
Free tmp (MB)	Available space in a temporary file system, in megabytes.
Max mem (MB)	Average total amount of maximum physical memory, in megabytes.
Max swap (MB)	Maximum available swap space, in megabytes.
Max tmp (MB)	Maximum available space in a temporary file system, in megabytes.
Used mem (MB)	Used physical memory, in megabytes.
Used swap (MB)	Used swap space, in megabytes.
CPU (%) used by jobs	Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.
Job used mem (MB)	Total of all physical memory used for jobs, in megabytes.
Job used swap (MB)	Total of all swap space used for jobs, in megabytes.

Cognos Cube Views

The Hourly Hardware cube includes the following Cognos cube views:

- Daily CPU utilization trends by host type
- Which hosts run too many jobs

30

The Hourly Hardware by Hostgroup Cube

The Hourly Hardware by Hostgroup cube contains hourly historical hardware data by host group. You can explore detailed hardware usage for all host groups by hour. This datamart is based on the Hourly Hardware by Hostgroup datamart. View data by year, quarter, month, week, day, or hour.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware by Hostgroup cube includes the following dimensions and child categories:

Host Group   Host group information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- ◆ [Host Group] | [Host]

Note: If a host belongs to multiple host groups, then the data is duplicated in each group.

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday	Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.

- ◆ Shift 1 | [Hour]
- ◆ Shift 2 | [Hour]
- ◆ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Measures

The Hourly Hardware by Hostgroup cube includes the following measures:

# of logins Number of logins.

CPU (%) Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

Used mem (%) Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).

Used swap (%) Used swap space as a percentage.

Used tmp (%) Used temp space as a percentage.

15s run queue Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

1min run queue Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

15min run queue Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Paging (pgs io/s) Paging activity, in pages in and out per second.

Disk IO (KB/second) Average disk I/O, for all disks on a host, in kilobytes per second.

Inter. idle time (min) Number of minutes since any interactive terminal activity.

Free mem (MB) Available physical memory, in megabytes.

Free swap (MB) Available swap space, in megabytes.

Free tmp (MB) Available space in a temporary file system, in megabytes.

Max mem (MB) Average total amount of maximum physical memory, in megabytes.

Max swap (MB) Maximum available swap space, in megabytes.

Max tmp (MB) Maximum available space in a temporary file system, in megabytes.

Used mem (MB) Used physical memory, in megabytes.

Used swap (MB) Used swap space, in megabytes.

CPU (%) used by jobs Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.

Job used mem (MB) Total of all physical memory used for jobs, in megabytes.

Job used swap (MB) Total of all swap space used for jobs, in megabytes.

Cognos Cube Views

The Hourly Hardware by Hostgroup cube includes the following Cognos cube views:

- Daily CPU utilization trends by host type
- Which hosts run too many jobs

31

The Hourly Hardware (with Boolean Resources) Cube

The Hourly Hardware (with Boolean Resources) cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Hourly Hardware (with Boolean Resources) datamarts. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, and host types by hour. For parallel jobs, the cube considers the first host in the execution host list as the execution host.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware (with Boolean Resources) cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

- Thursday
- Friday
- Saturday

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- Core hours
- Non-core hours

Shift Based on shift hours as defined in the console, under **Data Mappings > Shifts**.

- Shift 1 | [Hour]
- Shift 2 | [Hour]
- Shift 3 | [Hour]
- By Hour | [Hour]

Host Model Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Model] | [Host]

Host Type Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Type] | [Host]

Boolean Resource Host-based Boolean resources. When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled.

# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	◆ 1
	◆ 2

Measures

When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled. The Hourly Hardware (with Boolean Resources) cube includes the following measures:

# of logins Number of logins.

CPU (%) Average CPU usage as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.

CPU factor The relative CPU performance factor. The CPU factor is used to scale the CPU load value so that differences in CPU speeds are considered. The faster the CPU, the larger the CPU factor.

Used mem (%) Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).

Used swap (%) Used swap space as a percentage.

Used tmp (%) Used temp space as a percentage.

15s run queue Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.

1min run queue Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.

15min run queue Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.

Paging (pgs io/s) Paging activity, in pages in and out per second.

Disk IO (KB/second) Average disk I/O, for all disks on a host, in kilobytes per second.

Inter. idle time (min) Number of minutes since any interactive terminal activity.

Free mem (MB) Available physical memory, in megabytes.

Free swap (MB) Available swap space, in megabytes.

Free tmp (MB) Available space in a temporary file system, in megabytes.

Max mem (MB) Average total amount of maximum physical memory, in megabytes.

Max swap (MB) Maximum available swap space, in megabytes.

Max tmp (MB) Maximum available space in a temporary file system, in megabytes.

Used mem (MB) Used physical memory, in megabytes.

Used swap (MB) Used swap space, in megabytes.

CPU (%) used by jobs Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.

Job used mem (MB) Total of all physical memory used for jobs, in megabytes.

Job used swap (MB) Total of all swap space used for jobs, in megabytes.

Cognos Cube Views

The Hourly Hardware (with Boolean Resources) cube includes the following Cognos cube views:

- Daily CPU utilization trends by host type
- Which hosts run too many jobs

32

The Hourly Hardware (with Boolean Resources and Core Utilization) Cube

The Hourly Hardware (with Boolean Resources and Core Utilization) cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Hourly Hardware (with Boolean Resources and Core Utilization) and Hourly Core Utilization datamarts. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, and host types by hour. For parallel jobs, the cube considers the first host in the execution host list as the execution host.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware (with Boolean Resources and Core Utilization) cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Boolean Resource	Host-based Boolean resources. When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled.

# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	◆ 1
	◆ 2
Host-Core ID	Analytics-assigned ID that is unique to each core within each host.
	◆ [Host-Core ID]

Measures

When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled. The Hourly Hardware (with Boolean Resources and Core Utilization) cube includes the following measures:

# of logins	Number of logins.
CPU ut (%)	Average CPU utilization as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.
Used mem (%)	Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).
Used swap (%)	Used swap space as a percentage.
Used tmp (%)	Used temp space as a percentage.
15s run queue	Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.
1min run queue	Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.
15min run queue	Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.
Paging (pgs io/s)	Paging activity, in pages in and out per second.
Disk IO (KB/second)	Average disk I/O, for all disks on a host, in kilobytes per second.
Inter. idle time (min)	Number of minutes since any interactive terminal activity.
Free mem (MB)	Available physical memory, in megabytes.
Free swap (MB)	Available swap space, in megabytes.
Free tmp (MB)	Available space in a temporary file system, in megabytes.
Max mem (MB)	Average total amount of maximum physical memory, in megabytes.
Max swap (MB)	Maximum available swap space, in megabytes.
Max tmp (MB)	Maximum available space in a temporary file system, in megabytes.
Used mem (MB)	Used physical memory, in megabytes.
Used swap (MB)	Used swap space, in megabytes.
CPU (%) used by jobs	Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.
Job used mem (MB)	Total of all physical memory used for jobs, in megabytes.
Job used swap (MB)	Total of all swap space used for jobs, in megabytes.
Core ut (%)	Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.
# of hosts	Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.

Cognos Cube Views

The Hourly Hardware (with Boolean Resources and Core Utilization) cube includes the following Cognos cube views:

- Hourly CPU utilization trends by host type
- Which hosts run too many jobs

33

The Hourly Hardware (with Core Utilization) Cube

The Hourly Hardware (with Core Utilization) cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Hourly Hardware (with Core Utilization) and Hourly Core Utilization datamarts. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, and host types by hour. For parallel jobs, the cube considers the first host in the execution host list as the execution host.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Hardware (with Core Utilization) cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]

# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	◆ 1
	◆ 2
Host-Core ID	Analytics-assigned ID that is unique to each core within each host.
	◆ [Host-Core ID]

Measures

When measures are explored without the Boolean resource level, values for measures using sum as the regular rollup method are doubled. The Hourly Hardware (with Core Utilization) cube includes the following measures:

# of logins	Number of logins.
CPU ut (%)	Average CPU utilization as a percentage, based on all CPUs in the host. Note that the percentage is expressed as a whole number.
Used mem (%)	Used physical memory as a percentage. It is calculated as: (1 - **Free mem (MB)**/**Max mem (MB)**).
Used swap (%)	Used swap space as a percentage.
Used tmp (%)	Used temp space as a percentage.
15s run queue	Average number of processes that are ready to use the CPU during a 15-second interval. This is based on `lsload -l` in LSF.
1min run queue	Average number of processes that are ready to use the CPU during a one-minute interval. This is based on `lsload -l` in LSF.
15min run queue	Average number of processes that are ready to use the CPU during a 15-minute interval. This is based on `lsload -l` in LSF.
Paging (pgs io/s)	Paging activity, in pages in and out per second.
Disk IO (KB/second)	Average disk I/O, for all disks on a host, in kilobytes per second.
Inter. idle time (min)	Number of minutes since any interactive terminal activity.
Free mem (MB)	Available physical memory, in megabytes.
Free swap (MB)	Available swap space, in megabytes.
Free tmp (MB)	Available space in a temporary file system, in megabytes.
Max mem (MB)	Average total amount of maximum physical memory, in megabytes.
Max swap (MB)	Maximum available swap space, in megabytes.
Max tmp (MB)	Maximum available space in a temporary file system, in megabytes.
Used mem (MB)	Used physical memory, in megabytes.
Used swap (MB)	Used swap space, in megabytes.
CPU (%) used by jobs	Average CPU utilization by job, expressed as a percentage, based on all CPUs in the host for all hosts in the cluster.
Job used mem (MB)	Total of all physical memory used for jobs, in megabytes.
Job used swap (MB)	Total of all swap space used for jobs, in megabytes.
Core ut (%)	Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.
# of hosts	Average core utilization as a percentage, based on all cores in the host. Note that the percentage is expressed as a whole number.

Cognos Cube Views

The Hourly Hardware (with Core Utilization) cube includes the following Cognos cube views:

◆ Daily CPU utilization trends by host type
◆ Which hosts run too many jobs

34

The Hourly Host Availability Cube

The Hourly Host Availability cube contains hourly historical host down-time data. You can explore down-hours and down-times for all hosts, host models, and host types by hour. View data by year, quarter, month, week, day, or hour. This cube is based on the Hourly Host Availability and Host Properties datamarts.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Host Availability cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]

Measures

The Hourly Host Availability cube includes the following measures:

Down hours   Number of hours that a host was down during a specified time period.

# of down times   Number of times that a host went down during a specified time period.

> **Note**: There may be instances when the host is already down when the agent begins collecting host status data. If the host stays down during the whole data collection period (for example, one hour), than the down time will be reported as "0" (meaning it has not *gone* down during that time period).

Down percent   Down hours divided by a period of time.

Cognos Cube Views

The Hourly Host Availability Status cube does not include any Cognos cube views.

35

The Hourly Host Status Cube

The Hourly Host Status cube contains hourly historical host status data. You can explore the how long a host remains in a certain state for all hosts, host models, and host types by day. View data by year, quarter, month, week, day, or hour. This cube is based on the Hourly Host Status and Host Properties datamarts.

Contents
- "Dimensions" on page 218
- "Measures" on page 220
- "Cognos Cube Views" on page 221

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Host Status cube includes the following dimensions and child categories:

Status   Status Name

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆   [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆   [Yesterday]
- ◆   WTD | [Day]
- ◆   MTD | [Week] | [Day]
- ◆   QTD | [Month] | [Week] | [Day]
- ◆   YTD | |[Month] | [Week] | [Day]
- ◆   [Previous Week] | [Day]
- ◆   [Previous Month] | [Week] | [Day]
- ◆   [Previous Quarter] | [Month] | [Week] | [Day]
- ◆   [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆   Sunday
- ◆   Monday
- ◆   Tuesday
- ◆   Wednesday

- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour
Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
- ◆ Core hours
- ◆ Non-core hours

Shift
Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
- ◆ Shift 1 | [Hour]
- ◆ Shift 2 | [Hour]
- ◆ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Host Model
Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
- ◆ [Host Model] | [Host]

Host Type
Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
- ◆ [Host Type] | [Host]

Measures

The Hourly Host Status cube includes the following measure:

Duration   Length of time (in minutes) the host remained in a certain state during a specified time period.

Cognos Cube Views

The Hourly Host Status cube does not include any Cognos cube views.

36

The Hourly Job Resource Consumption Cube

The Hourly Job Resource Consumption cube lets you explore historical data about resources reserved and consumed by jobs. A job resource is the LSF resource used in the resource string when submitting LSF jobs with the -R option, specifically in the rusage [ ] section. It is based on data collected from the lsb.acct file, collected in the Hourly Job Reserved Resource Usage datamart.

Dimensions

Abbreviations

WTD Week to Date

MTD Month to Date

QTD Quarter to Date

YTD Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Job resource A resource that can be reserved. Even if `decay` or `duration` is used in the resource string, the entire run time is still added to the job.

> **Note:** To track how a license is reserved, you must specify it in the `rusage` section.

Job resource consumption Job resource consumption equals the total number of minutes that all resources are reserved.

Dimensions

The Hourly Job Resource Consumption cube includes the following dimensions and child categories:

Resource ◆ [All Reserved Resources] | [Reserved Resource]

Cluster ◆ Cluster | [Host]

Date The date when the job finished running, based on the date and time according to the LSF master host.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆ [Yesterday]
◆ WTD | [Day]
◆ MTD | [Week] | [Day]
◆ QTD | [Month] | [Week] | [Day]
◆ YTD | |[Month] | [Week] | [Day]

◆ [Previous Week] | [Day]

◆ [Previous Month] | [Week] | [Day]

◆ [Previous Quarter] | [Month] | [Week] | [Day]

◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

◆ Sunday

◆ Monday

◆ Tuesday

◆ Wednesday

◆ Thursday

◆ Friday

◆ Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours

◆ Non-core hours

Shift   Based on shift hours as defined in the console, under **Data Mappings > Shifts**.

◆ Shift 1 | [Hour]

◆ Shift 2 | [Hour]

◆ Shift 3 | [Hour]

◆ By Hour | [Hour]

Project   ◆ [User-defined Project]

User   ◆ [All Users] | [User]

Measures

The Hourly Job Resource Consumption cube includes the following measures:

Usage (minute)   The total number of minutes reserved by all job resources. By default, this shows the amount of time in minutes that each resource is reserved. (Resource names are rows, and cluster names are columns.)

Cost   Cost is calculated as **Usage (minute)** x **Resource Rate**. The Resource Rate is defined in the console under **Clusters > [*Cluster Name*] > Resource Costs > Job Cost Rate**. You can specify the job cost rate as well as the start date for applying this rate.

Cognos Cube Views

The Hourly Job Resource Consumption cube includes the following Cognos cube views:

- ◆ Resource consumption by business hour
- ◆ Resource consumption by project
- ◆ Resource consumption by user
- ◆ Resource consumption by weekday

37

The Hourly Job Slot Usage Cube

The Hourly Job Slot Usage cube lets you explore historical hourly job slot utilization and usage data according to group type. It is based on the Hourly Job Slot Usage by Group datamart. You can explore job slot utilization and usage information by user, host, project, queue, job group, and job name.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Job Slot Usage cube includes the following dimensions and child categories:

Group Type
- By Cluster
- By Host
- By Jobgroup
- By Jobname
- By Project
- By Queue
- By User
- By Usergroup

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]

- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Shift
- ◆ By Shift
  - ❖ Shift 1 | [Hour]
  - ❖ Shift 2 | [Hour]
  - ❖ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
- ◆ Core hours
- ◆ Non-core hours

Weekday Based on the date and time according to the LSF master host.
- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Measures

The Hourly Job Slot Usage cube includes the following measures:

Max Slot Usage Top number of slots used at one time for each group type; measures the number of job slots used for a group type.

Min Slot Usage Minimum number of slots used at one time for each group type; measures the number of job slots used for a group type.

Avg Slot Usage Average number of slots used at one time for each group type; measures the number of job slots used for a group type.

Max Slot Utilization (%) The maximum slot utilization for each group type, expressed as a percentage.

Min Slot Utilization (%) The minimum slot utilization for each group type, expressed as a percentage.

Avg Slot Utilization (%) The average slot utilization for each group type, expressed as a percentage.

Cognos Cube Views

The Hourly Job Slot Usage cube does not include any Cognos cube views.

38

The Hourly Job Statistics Cube

The Hourly Job Statistics cube lets you explore historical job data about jobs in the cluster in any state (running, pending, and suspended). It is based on the Hourly Job Statistics datamart. You can explore the job statistics for all hosts, projects, queues, users, and clusters.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Notes about the Group, Job Type, and Job Status dimensions

In order to obtain meaningful data, you must specify child categories (i.e. filter) all of the following dimensions:

◆ Group
◆ Job Type
◆ Job Status

You can filter these dimensions from the Dimension line. For additional information, see "Filtering data" on page 237.

Dimensions

The Hourly Job Statistics cube includes the following dimensions and child categories:

Group   To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 237.

◆ By Host
◆ By Project
◆ By Queue
◆ By User
◆ By Cluster

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]

- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Shift Based on shift hours as defined in the console, under **Data Mappings > Shifts**.

- ◆ Shift 1 | [Hour]
- ◆ Shift 2 | [Hour]
- ◆ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Job Type To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 237.

- ◆ Batch
- ◆ Interactive
- ◆ Parallel

Note: Parallel jobs are displayed as a string of host names.

Job Status To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 237.

- ◆ Pending
- ◆ Running

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Filtering data

In order to obtain meaningful results, you must filter the following dimensions to child categories:

- ◆ Group

- Job Type
- Job Status

Otherwise, the data is averaged together for all the child categories of each dimension, which is not meaningful.

For example, suppose that in the Group dimension, you want to organize jobs statistics by project. You can do one of the following:

- Click the **Project** hypertext link in the row heading.
- In the Dimension line, click **Group > Project**.

Measures

The Hourly Job Statistics cube includes the following measures:

Avg Job Number — The average number of jobs, based on a 15-minute default collection interval that is aggregated hourly. This is the average, based on weight by the sampling count measure (below). Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Min Job Number — The minimum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. This is the minimum, during hourly aggregation, of all the minimum values collected during each collection period. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Max Job Number — The maximum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. This is the maximum, during hourly aggregation, of all the maximum values collected during each collection period. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

How the Summary is reached

You may notice that the Total Avg number of jobs in the Summary line of the Hourly Job Statistics report does not match the what is stated in the Summary line of the Hourly Job Statistics cube. Both calculations are correct; however, the methods used to calculate an average are different.

Due to a different drill-down and roll-up for a cube's multiple dimensions, a weighted average is used (divide by weight average), while the report summary uses the normal average (divide by number of jobs).

Cognos Cube Views

The Hourly Job Statistics cube includes the following Cognos cube views:

- Number of pending jobs per cluster by business hour
- Number of running jobs per cluster by business hour
- Number of jobs per queue by job status

39

The Hourly Job Throughput Cube

The Hourly Job Throughput cube lets you explore the job throughput of different users and projects. The cube contains historical job throughput data. It is based on data collected from the `lsb.acct` file, collected in the Host Properties and Hourly Job Throughput datamarts. The dynamic data (from the Hourly Job Throughput datamart) is collected hourly and aggregated daily. You can explore the job throughput for all queues, projects, users, and hosts.

Contents

Dimensions

Abbreviations

WTD Week to Date

MTD Month to Date

QTD Quarter to Date

YTD Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Job Throughput cube includes the following dimensions and child categories:

Cluster Cluster | [Host]

Date The date when the job finished running, based on the date and time according to the LSF master host.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Queue	◆ [Priority] | [Level]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Project	◆ [User-defined Project]
User	◆ [User] | [User]
# of CPUs per Host	The number of physical processors in the host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	The number of cores in each physical processor in the host.
	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	The maximum number of threads in each core in the host.
	◆ 1
	◆ 2

Measures

Data collection

Data about job throughput is collected in "real time", so that it does not include data about running or pending jobs. Therefore, jobs that were submitted before midnight but are still running or pending at the last midnight before the cube build are not included in the count of submitted jobs. Consequently, the data in the cubes data changes frequently, especially when there are lots of long jobs, which run for more than one day.

Consider the example of a single cluster that is monitored by Analytics with the following average hourly statistics:

◆ 300 jobs are submitted
◆ 200 jobs finish running with done status
◆ 50 jobs finish running with exit status
◆ 20 jobs never run
◆ 30 jobs are still running when Analytics analyzes the data

The average number of submitted jobs in the resulting cube will increase from the time Analytics originally analyzes the data to once the 30 running jobs finish running.

Calculation of measures

The measures in this cube are the total number of jobs during the specified period. You can drill into the data to see the hourly totals.

The Job Throughput cube includes the following measures:

of submitted jobs

Total number of submitted jobs.

of done jobs

Total number of done jobs.

of exited jobs

Total number of exited jobs.

Cognos Cube Views

The Hourly Job Throughput cube includes the following Cognos cube views:

◆ Overall job throughput by business hour
◆ Job throughput by host model
◆ Job throughput by host type
◆ Job throughput by queue and business hour
◆ Number of submitted jobs by user and business hour

40

The Hourly Jobs by Pending Reason Cube

The Hourly Jobs by Pending Reason cube lets you explore the reasons jobs are pending. The cube contains historical data about the reasons jobs are pending. It is based on the Hourly Jobs by Pending Reason datamart. You can explore these jobs by pending reason, type of pending reason, and cluster.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Other notes about the Group dimension

In order to obtain meaningful data, you must specify child categories (i.e. filter) for the Group dimension. You can filter this dimension from the Dimension line. For additional information, see "Filtering data" on page 249.

Dimensions

The Jobs By Pending Reason cube includes the following dimensions and child categories:

Group   To obtain meaningful data, you must specify one of the child categories. For additional information, see "Filtering data" on page 249.

- ◆ By Pending Reason
  - ❖ Any pending reason in LSF
- ◆ By Pending Reason Type
  - ❖ Any pending reason type in LSF, such as User-Related Reasons, Job-Related Reasons, and Host-Related Reasons
- ◆ By Cluster

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]

- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | | [Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Shift Based on shift hours as defined in the console, under **Data Mappings > Shifts**.

- ◆ Shift 1 | [Hour]
- ◆ Shift 2 | [Hour]
- ◆ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Filtering data

In order to obtain meaningful results, you must filter the Group dimension to a child category. In other words, you need to specify whether to display the data by project, queue, user, pending reason, or cluster. Otherwise, the data is averaged for all these groups together, which is not meaningful.

For example, suppose that in the Group dimension, you want to organize jobs statistics by project. You can do one of the following:

- ◆ Click the **By Project** hypertext link in the row heading.
- ◆ In the Dimension line, click **Group > By Project**.

Measures

The Hourly Jobs By Pending Reason cube includes the following measures:

Avg Job Number The average number of jobs, based on a 15-minute default collection interval that is aggregated hourly. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Min Job Number The minimum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Max Job Number The maximum number of jobs, based on a 15-minute default collection interval that is aggregated hourly. Note that you can modify the default collection interval when you deploy the agent. For additional information, see *Administering Platform Analytics*.

Cognos Cube Views

The Hourly Jobs By Pending Reason cube includes the following Cognos cube view report:

◆ Pending jobs per cluster by pending reason

41

The Hourly License Consumption Cube

The Hourly License Consumption cube lets you explore up to two weeks of hourly license consumption and utilization, to ensure that the license service level is good enough in your cluster. It contains information about license vendors, license servers, user groups, and execution host groups. It is based on the Hourly License Consumption datamart.

◆ This cube samples data frequently.

◆ Only recent data is available; data expires after 30 days.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Consumption cube includes the following dimensions and child categories:

License vendor
: ◆ All vendors | [License vendor] | [License feature]

Date
: The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday
: Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users| [User]
Host	◆ All Hosts | [Host]
License server	◆ [License Server]

Measures

The Hourly License Consumption cube includes the following measures:

License minutes (number x minutes)

The total number of minutes that all licenses are used.

Cost

The dollar cost of a specific to a user, project, etc. Cost is calculated as **LIcense minutes** x **License Rate**. You can enter the license rate in the console under **Resource Costs > License Rate**. You can specify the license rate as well as the start date for applying this rate.

Average # of used licenses

The average number of used licenses. It is calculated as: **License minutes/# of hours**/60 (minutes).

The average number of used licenses.

Cognos Cube Views

The Hourly License Consumption cube does not include any Cognos cube views.

42

The Hourly License Denial Cube

The Hourly License Denial cube lets you explore the license denial and utilization, and to ensure that the license service level is good enough in your cluster. It contains information about license vendors, license servers, user groups, and execution host groups. It is based on the Hourly License Denial datamart.

The data in the Hourly License Denial cube is based on the FLEXlm debug log files. If you disable the creation of these debug log files in the Options file, the License Denial cube will be empty. For additional information about the FLEXlm debug log files, refer to the *FLEXlm End Users Guide*.

◆ This cube samples data frequently.

◆ Only recent data is available; data expires after 30 days.

◆ The data in this cube is aggregated hourly.

Differences between the License Denial cubes and the License Denial by Feature cubes

License Denial cubes	License Denial by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers
◆ cannot identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is identified	◆ can identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is recognized as false and ignored

Configuration   In order for the license server name to be displayed in license denial cubes and Cognos cube views, you must configure the license log file in the following format:

```
Fullpath/license_log_file | License_server_name
```

Contents
◆ "Dimensions" on page 260
◆ "Measures" on page 262
◆ "Cognos Cube Views" on page 263

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Denial cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users | [User]
Host	◆ All Hosts | [Host]
License server	◆ [License Server]

Measures

The Hourly License Denial cube includes the following measures:

License denials

The total number of license denials.

Cognos Cube Views

The License Denial cube includes the following Cognos cube views:

- Features with the most denials
- Hosts with the most denials
- Servers with the most denials
- Users with the most denials

43

The Hourly License Denial By Feature Cube

The Hourly License Denial By Feature cube lets you explore the historical data about global-wide license denial by feature and the service level of license servers. It contains information about license vendors, license servers, users, and execution hosts. It is based on the Hourly License Denial By Feature datamart.

The data in the Hourly License Denial By Feature cube is based on the FLEXlm debug log files. If you disable the creation of these debug log files in the Options file, the License Denial cube will be empty. For additional information about the FLEXlm debug log files, refer to the *FLEXlm End Users Guide.*

◆ This cube samples data frequently.

◆ Only recent data is available; data expires after 30 days.

◆ The data in this cube is aggregated hourly.

Differences between the License Denial cubes and the License Denial by Feature cubes

License Denial cubes	License Denial by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers
◆ cannot identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is identified	◆ can identify false denials across license servers, for example: if an application is denied a license by one server but checks out a license from another server, the denial by the first server is recognized as false and ignored

Configuration   In order for the license server name to be displayed in license denial cubes and Cognos cube views, you must configure the license log file in the following format:

```
Fullpath/license_log_file | License_server_name
```

Contents
◆ "Dimensions" on page 266

◆ "Measures" on page 268

◆ "Cognos Cube Views" on page 269

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Denial By Feature cube includes the following dimensions and child categories:

License vendor
: ◆ All vendors | [License vendor] | [License feature]

Date
: The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday
: Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours
◆ Non-core hours

User	◆ All Users| [User]
Host	◆ All Hosts | [Host]

Measures

The Hourly License Denial By Feature cube includes the following measures:

License denials

The total number of license denials.

Cognos Cube Views

The Hourly License Denial By Feature cube does not include any Cognos cube views.

44

The Hourly License Usage By User Cube

The Hourly License Usage By User cube lets you explore two weeks of data about hourly license usage from a user perspective. It contains information about license vendors, license servers, user groups, and execution host groups. It is based on the Hourly License Usage datamart. For information about license utilization from a license server perspective, see "The Hourly License Utilization Cube" on page 295.

◆ This cube samples data frequently.

◆ Only recent data is available; data expires after 30 days.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Usage By User cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

User
- All Users | [User]

Date The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | | [Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday

◆ Friday

◆ Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours

◆ Non-core hours

License server   ◆ All License Servers | [License Server]

Measures

The Hourly License Usage By User cube includes the following measures:

Average usage by user

The average license usage by user. Usage is expressed as a number.

Peak usage by user

The maximum license usage by user. Usage is expressed as a number.

Average # of available licenses

The average total number of available licenses in a specified time period.

Cognos Cube Views

The Hourly License Usage By User cube includes the following Cognos cube view report:

- License usage and available licenses by business hour

45

The Hourly License Usage By User Group Cube

The Hourly License Usage By User Group cube lets you explore two weeks of data about hourly license usage from a user group perspective. It contains information about license vendors, license servers, and user groups. It is based on the Hourly License Usage by User Group datamart. For information about license utilization from a license server perspective, see "The Hourly License Utilization Cube" on page 295.

Note: By default, Analytics does not build this cube. For information about building cubes, refer to *Administering Platform Analytics*.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Usage By User Group cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Usergroup
- All Usergroups| [Usergroup]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday

◆ Friday

◆ Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours

◆ Non-core hours

License server   All License Servers | [License Server]

Measures

The Hourly License Usage By User Group cube includes the following measures:

Average usage by usergroup

The average license usage by user group. Usage is expressed as a number.

Peak usage by usergroup

The maximum license usage by user group. Usage is expressed as a number.

Average # of available licenses

The average total number of available licenses in a specified time period.

Cognos Cube Views

The Hourly License Usage By User Group cube does not include any Cognos cube views.

46

The Hourly License Usage Statistics by Feature Cube

The Hourly License Usage Statistics by Feature cube lets you explore two weeks of hourly license usage statistics by feature. It contains information about license denial, utililization, and consumption, all by feature in one cube. It is based on the Hourly License Usage Statistics by Feature - Denial, Hourly License Usage Statistics by Feature - Utilization, and Hourly License Usage Statistics by Feature - Consumption datamarts.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Usage Statistics by Feature cube includes the following dimensions and child categories:

License vendor

- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users| [User]
Host	◆ All Hosts | [Host]

Measures

The Hourly License Usage Statistics by Feature cube includes the following measures:

License denials

The total number of license denials.

Average usage by feature

The average license usage by feature. Usage is expressed as a number.

Peak usage by feature

The maximum license usage by feature. Usage is expressed as a number.

Average utilization by feature(%)

The average license utilization for each feature over a period of time, expressed as a percentage.

Average # of available licenses

The average total number of available licenses in a specified time period.

Peak # of available licenses by feature

The peak number of licenses available for each feature over a period of time.

Peak utilization by feature (%)

The peak license utilization for each feature over a period of time, expressed as a percentage.

Utilization at peak usage by feature (%)

The peak license utilization for each feature at the time when the maximum number of licenses are being used.

# of reserved licenses by feature

The number of reserved licenses for each feature.

License available minutes

The amount of time that licenses were available, in minutes. This is calculated as **number of licenses_available × minutes available**.

License minutes

The amount of time that licenses were used, in minutes. This is calculated as **number of licenses_used × minutes used**.

Cost

Cost is calculated as **minutes_used × Resource Rate**. The Resource Rate is defined in the console under **Clusters > [*Cluster Name*] > Resource Costs > License Cost Rate**. You can specify the job cost rate as well as the start date for applying this rate.

Cognos Cube Views

The Hourly License Usage Statistics by Feature cube does not include any Cognos cube views.

47

The Hourly License Usage Statistics Cube

The Hourly License Usage Statistics cube lets you explore two weeks of hourly license usage statistics. It contains information about license denial, utililization, and consumption in one cube. It is based on the Hourly License Usage Statistics - Denial, Hourly License Usage Statistics - Utilization, and Hourly License Usage Statistics - Consumption datamarts.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Usage Statistics cube includes the following dimensions and child categories:

License vendor

- ◆ All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
User	◆ All Users| [User]
Host	◆ All Hosts	[Host]
License server	◆ All License Servers	[License Server]

Measures

The Hourly License Usage Statistics cube includes the following measures:

License denials

The total number of license denials.

Average usage by server

The average license usage by server. Usage is expressed as a number.

Peak usage by server

The maximum license usage by server. Usage is expressed as a number.

Average utilization by server (%)

The average license utilization on each license server over a period of time, expressed as a percentage.

Peak utilization by server (%)

The peak license utilization on each license server over a period of time, expressed as a percentage.

Average # of available licenses

The average total number of available licenses in a specified time period.

Utilization at peak usage by server (%)

The peak license utilization on each license server at the time when the maximum number of licenses are being used.

Peak # of available licenses by server

The peak number of licenses available on each license server over a period of time.

of reserved licenses by server

The number of reserved licenses on each license server.

License available minutes

The amount of time that licenses were available, in minutes. This is calculated as **number of licenses_available** $\times$ **minutes available**.

License minutes

The amount of time that licenses were used, in minutes. This is calculated as **number of licenses_used** $\times$ **minutes used**.

Cost

Cost is calculated as **minutes_used** $\times$ **Resource Rate**. The Resource Rate is defined in the console under **Clusters > [*Cluster Name*] > Resource Costs > License Cost Rate**. You can specify the job cost rate as well as the start date for applying this rate.

Cognos Cube Views

The Hourly License Usage Statistics cube does not include any Cognos cube views.

48

The Hourly License Utilization Cube

The Hourly License Utilization cube lets you explore the license utilization from the license server perspective. It contains historical license utilization data, and it is based on the Hourly License Utilization datamart. It does not contain information about individual users. For information about license usage from a user perspective, see "The Hourly License Usage By User Cube" on page 271.

Differences between the License Utilization cubes and the License Utilization by Feature cubes

License Utilization cubes	License Utilization by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Utilization cube includes the following dimensions and child categories:

License vendor
- ◆ All vendors | [License vendor] | [License feature]

Date
: The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

License server
- ◆ [License Server]

Measures

The Hourly License Utilization cube includes the following measures:

Average usage by server

The average number of licenses in use on each license server.

Peak usage by server

The maximum number of licenses in use on each license server.

Average utilization by server(%)

The average license utilization on each license server over a period of time, expressed as a percentage.

Peak utilization by server (%)

The peak license utilization on each license server over a period of time, expressed as a percentage.

Average # of available licenses

The average number of licenses available on the license server.

Utilization at peak usage by server (%)

The utilization for each feature during peak usage, expressed as a percentage.

Peak # of available licenses by server

The peak total number of licenses available on each license server over a period of time.

of reserved licenses by server

The number of licenses reserved on each license server.

Cognos Cube Views

The Hourly License Utilization cube does not include any Cognos cube views.

49

The Hourly License Utilization By Feature Cube

The Hourly License Utilization By Feature cube lets you explore the license utilization from the individual feature perspective. It contains historical hourly data about global-wide license utilization by feature, and it is based on the Hourly License Utilization By Feature datamart. It does not contain information about individual users. For information about license usage from a user perspective, see "The Hourly License Usage By User Cube" on page 271.

Differences between the License Utilization cubes and the License Utilization by Feature cubes

License Utilization cubes	License Utilization by Feature cubes
◆ include a Server dimension	◆ no Server dimension
◆ reflects service level by license server	◆ reflects the service level of a license feature across all license servers

Contents
- ◆ "Dimensions" on page 300
- ◆ "Measures" on page 301
- ◆ "Cognos Cube Views" on page 302

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly License Utilization By Feature cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Measures

The Hourly License Utilization By Feature cube includes the following measures:

Average usage by feature

The average number of licenses in use, by feature.

Peak usage by feature

The maximum number of licenses in use, by feature.

Average utilization by feature (%)

The average license utilization for each feature over a period of time, expressed as a percentage.

Average # of available licenses

The total number of licenses available on the license server.

Peak # of available licenses by feature

The peak total number of licenses available on the license server over a period of time, by feature.

Peak utilization by feature (%)

The peak license utilization for each feature over a period of time, expressed as a percentage.

Utilization at peak usage by feature (%)

The utilization on each license server during peak usage, by feature, expressed as a percentage.

of reserved licenses by feature

The number of licenses reserved for each feature.

Cognos Cube Views

The Hourly License Utilization By Feature cube does not include any Cognos cube views.

50

The Hourly Shared Resource Cube

The Hourly Shared Resource cube contains historical data about the percentage of utilization and usage of shared resources. It shows only numeric, dynamic, user-defined resources in LSF. Information about shared resources is based on data collected from LSF, in the Hourly Shared Resource Usage datamart.

◆ The data in this cube is aggregated hourly.
◆ The data in this cube expires after 30 days.

Contents ◆
◆
◆

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Shared resource
: A cluster-based LSF resource that is shared among hosts, e.g. a license.

Cluster
: A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Shared Resource Usage cube includes the following dimensions and child categories:

Resource
: ◆ [All Shared Resources] | [Shared Resource]

Cluster
: ◆ Cluster | [Host]

Date
: The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆ [Yesterday]
◆ WTD | [Day]
◆ MTD | [Week] | [Day]
◆ QTD | [Month] | [Week] | [Day]
◆ YTD | |[Month] | [Week] | [Day]
◆ [Previous Week] | [Day]
◆ [Previous Month] | [Week] | [Day]
◆ [Previous Quarter] | [Month] | [Week] | [Day]
◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Business Hour
: Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

◆ Core hours

◆ Non-core hours

Shift Based on core hours as defined in the console, under **Data Mappings > Shifts**.

◆ Shift 1 | [Hour]
◆ Shift 2 | [Hour]
◆ Shift 3 | [Hour]
◆ By Hour | [Hour]

Measures

In order to obtain meaningful data for the **Average usage**, **Average utilization (%)**, and **Average total available** measures (described below), you must make sure that there is data in the LIC_FEATURE_NAME and TOTAL_AVAIL columns in the WI_LICENSETOTAL table. In general, this table contains data from the LICENSEUSAGE table. However, if the you disabled this table, or if the shared resources are not licenses, you need enter the following information in the WI_LICENSETOTAL table:

◆ shared resource name in the LIC_FEATURE_NAME column
◆ total number of shared resources in the TOTAL_AVAIL column

You can enter this information in one of the following ways:

◆ use a database tool
◆ use a script, such as SQL

Average usage

The average number of shared resources.

Average utilization (%)

The average utilization of shared resources by server, expressed as a percentage.

Average total available

The average total number of available shared resource, including licenses in use and licenses that are free for use.

Average free

The average total number of licenses that are free for use.

Cognos Cube Views

The Hourly Shared Resource cube does not include any Cognos cube views.

51

The Hourly Workload Cube

The Hourly Workload cube lets you explore how much hardware and time different users and projects consume, and discover whether cluster usage actually maps to your business goals and strategies. It contains historical workload data about jobs, queues, projects, and hosts. It is based on data collected from the `lsb.acct` file, collected in the Host Properties and Host Job Information datamarts. You can explore the workload for all projects, queues, users, and hosts.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly Workload cube includes the following dimensions and child categories:

Cluster   ◆   Cluster | [Host]

Date   The date when the job finished running, based on the date and time according to the LSF master host.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆   [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆   [Yesterday]
◆   WTD | [Day]
◆   MTD | [Week] | [Day]
◆   QTD | [Month] | [Week] | [Day]
◆   YTD | |[Month] | [Week] | [Day]
◆   [Previous Week] | [Day]
◆   [Previous Month] | [Week] | [Day]
◆   [Previous Quarter] | [Month] | [Week] | [Day]
◆   [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   The weekday when the job was finished.

◆   Sunday
◆   Monday
◆   Tuesday
◆   Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on core hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Date (Submit)	The date when the job was submitted, based on the date and time according to the LSF master host.
	◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
	◆ QTD | [Month] | [Week] | [Day]
	◆ Prior QTD | [Month] | [Week] | [Day]
	◆ QTD Grouped
	❖ Prior QTD | [Month] | [Week] | [Day]
	❖ QTD | [Month] | [Week] | [Day]
	❖ QTD Change
	❖ QTD Growth
	◆ YTD | [Quarter] | [Month] | [Week] | [Day]
	◆ YTD Grouped
	❖ YTD | [Quarter] | [Month] | [Week] | [Day]
	◆ MTD | [Week] | [Day]
	◆ Prior MTD | [Week] | [Day]
	◆ MTD Grouped
	❖ Prior MTD | [Week] | [Day]
	❖ MTD | [Week] | [Day]
	❖ MTD Change
	❖ MTD Growth
	◆ WTD | [Day]
	◆ Prior WTD | [Day]
	◆ WTD Grouped
	❖ Prior WTD | [Day]
	❖ WTD | [Day]
	❖ WTD Change
	❖ WTD Growth
Weekday (Submit)	The weekday when the job was submitted, based on the date and time according to the LSF master host.
	◆ Sunday

- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour (Submit) The type of business hour when the job was submitted, based on the date and time according to the LSF master host.

- Core hours
- Non-core hours

Note: Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

Date/Start The date when the job started running, based on the date and time according to the LSF master host.

- No start date
- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- Prior QTD | [Month] | [Week] | [Day]
- QTD Grouped
  - Prior QTD | [Month] | [Week] | [Day]
  - QTD | [Month] | [Week] | [Day]
  - QTD Change
  - QTD Growth
- YTD | [Quarter] | [Month] | [Week] | [Day]
- YTD Grouped
  - YTD | [Quarter] | [Month] | [Week] | [Day]
- MTD | [Week] | [Day]
- Prior MTD | [Week] | [Day]
- MTD Grouped
  - Prior MTD | [Week] | [Day]
  - MTD | [Week] | [Day]
  - MTD Change
  - MTD Growth
- WTD | [Day]
- Prior WTD | [Day]
- WTD Grouped
  - Prior WTD | [Day]
  - WTD | [Day]
  - WTD Change
  - WTD Growth

Weekday/Start The weekday when the job started running, based on the date and time according to the LSF master host.

	◆ Sunday
	◆ Monday
	◆ Tuesday
	◆ Wednesday
	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour/Start	The type of business hour when the job started running, based on the date and time according to the LSF master host.
	◆ Core hours
	◆ Non-core hours

Note: Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

Project	◆ [User-defined Project]
User	◆ [All Users] | [User]
Queue	◆ [Priority] | [Level]
Application Tag	◆ [Application Tag]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Execution Host Type] | [Host]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Execution Host Model] | [Host]
# of CPUs per Host	The number of physical processors in the execution host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	◆ 1
	◆ 2
Host Type (Submit)	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Submission Host Type] | [Host]
Host Model (Submit)	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Submission Host Model] | [Host]
Job Type	◆ [User-defined job]

Job Group	Name of the job group as defined in Platform LSF.
	◆ [Job Group]
# of CPUs per Job	The number of CPUs used to run the job. This is not necessarily the same as the number of CPUs on the host.
	◆ 0.0 (The job was never executed.)
	◆ 1.0
	◆ 4.0
Job Exit Status	◆ DONE
	◆ EXIT
Job Exit Code	◆ [Job Exit Code]
Pending Time Rank	◆ 0 sec to < 5 sec
	◆ 5 sec to < 30 sec
	◆ 30 sec to < 5 min
	◆ 5 min to < 60 min
	◆ 60 min to < 1 day
	◆ One day or more
Run Time Rank	◆ 0 sec to < 1 min
	◆ 1 min to < 5 min
	◆ 5 min to < 20 min
	◆ 20 min to < 60 min
	◆ 60 min to < 4 hour
	◆ 4 hour to < 1 day
	◆ One day or more

CPU Time Rank
- 0 sec to < 1 min
- 1 min to < 5 min
- 5 min to < 20 min
- 20 min to < 60 min
- 60 min to < 4 hour
- 4 hour to < 1 day
- One day or more

Memory Usage Rank
- 0 GB to < 1 GB
- 1 GB to < 2 GB
- 2 GB to < 4 GB
- 4 GB to < 8 GB
- 8 GB to < 16 GB
- 16 GB to < 32 GB
- 32 GB or more

User Group   Name of user group as defined in Platform LSF.

- [User Group]

Note: If a user belongs to multiple user groups, then the job submitted by that user is not associated with a group unless "-G"' is used to indicate the group when submitting the job.

Job ID   The Job ID as assigned in Platform LSF.

- [Job ID]

Note: To use this dimension, you must enable the "Include Job ID in Workload Cubes" feature in the Platform Analytics Console (click **Addons** in the navigation tree).

Measures

The Hourly Workload cube includes the following measures:

of jobs

Total number of jobs.

Job run time (min)

Total run time of all jobs, in minutes.

Job pend time (min)

Total pend time of all jobs, in minutes.

Job turnaround time (min)

Total time elapsed between job submission and job finish.

CPU time (sec)

Total CPU time used by all jobs, in seconds.

Used mem (MB)

Total memory used by all jobs, in megabytes.

Used swap (MB)

Total swap space used by all jobs, in megabytes.

CPU time/Job run time (%)

Total real CPU time used by a job, divided by the total time elapsed between job start and job finish.

Avg run time/job (min)

Average run time of a job, in minutes.

Avg pend time/job (min)

Average pend time of a job, in minutes.

Avg turnaround time/job (min)

Average time elapsed between job submission and job finish for all specified jobs.

Job run time/Turnaround time (%)

Total time elapsed between job start and job finish, divided by the total time elapsed between job submission and job finish.

Avg CPU time/job (min)

Average CPU time used by a job, in minutes.

Avg used mem/job (MB)

Average memory used by a job, in megabytes.

Avg used swap/job (MB)

Average swap space used by a job, in megabytes.

Cognos Cube Views

The Hourly Workload cube includes the following Cognos cube views:

- Consumed resources by job exist status and cluster
- Jobs ranked by pending time
- Throughput efficiency by cluster and date
- Types of hosts completing jobs the fastest
- Projects consuming the most CPU time
- Projects running the most jobs
- Projects waiting the longest for cluster resources
- Projects wasting the most CPU time
- Projects with the most CPU-intensive jobs
- Throughput efficiency by cluster and project
- Jobs in queue ranked by pending time
- Number of done and exited jobs by queue
- Throughput efficiency by queue
- Users consuming the most CPU time
- Users waiting the longest to execute jobs
- Users with the most jobs

52

The Hourly Workload License Cube

The Hourly Workload License cube lets you explore the hourly historical workload license data obtained from your LSF cluster and the FLEXnet Manager server. This cube is based on the Hourly Workload License datamart.

Note: Macrovision's FLEXnet Manager is a central management tool for all FLEXlm servers in an enterprise environment. In order to build and explore the data that makes up this cube, you require a FLEXnet Manager installation.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The Hourly Workload License cube includes the following dimensions and child categories:

Cluster   Cluster | [Name]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Shift   Based on shift hours as defined in the console, under **Data Mappings > Shifts**.

- ◆ Shift 1 | [Hour]
- ◆ Shift 2 | [Hour]
- ◆ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Date (Submit)   The date when the job was submitted of which you want sampled, based on the date and time according to the LSF master host.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ Current Day
- ◆ Yesterday
- ◆ WTD | [Day]
- ◆ Prior WTD | [Day]
- ◆ WTD Grouped
    - ❖ Prior WTD | [Day]
    - ❖ WTD | [Day]
    - ❖ WTD Change
    - ❖ WTD Growth
- ◆ MTD | [Week] | [Day]
- ◆ Prior MTD | [Week] | [Day]
- ◆ MTD Grouped
    - ❖ Prior MTD | [Week] | [Day]
    - ❖ MTD | [Week] | [Day]
    - ❖ MTD Change
    - ❖ MTD Growth
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ Prior QTD | |[Month] | [Week] | [Day]
- ◆ QTD Grouped
    - ❖ Prior QTD | [Month] | [Week] | [Day]
    - ❖ QTD | [Month] | [Week] | [Day]
    - ❖ QTD Change
    - ❖ QTD Growth
- ◆ YTD | [Quarter] | [Week] | [Day]
- ◆ Prior YTD | [Quarter] | [Week] | [Day]
- ◆ YTD Grouped
    - ❖ Prior YTD | [Quarter] | [Week] | [Day]
    - ❖ YTD | [Quarter] | [Week] | [Day]
    - ❖ YTD Change
    - ❖ YTD Growth

Weekday (Submit) The weekday when the job was submitted of which you want sampled, based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour (Submit) The type of business hour when the job was submitted, based on the date and time according to the LSF master host.

- Core hours
- Non-core hours

Note: Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

Date/Start The date when the job started running of which you want sampled, based on the date and time according to the LSF master host.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- Current Day
- Yesterday
- WTD | [Day]
- Prior WTD | [Day]
- WTD Grouped
  - Prior WTD | [Day]
  - WTD | [Day]
  - WTD Change
  - WTD Growth
- MTD | [Week] | [Day]
- Prior MTD | [Week] | [Day]
- MTD Grouped
  - Prior MTD | [Week] | [Day]
  - MTD | [Week] | [Day]
  - MTD Change
  - MTD Growth
- QTD | [Month] | [Week] | [Day]
- Prior QTD | |[Month] | [Week] | [Day]
- QTD Grouped
  - Prior QTD | [Month] | [Week] | [Day]
  - QTD | [Month] | [Week] | [Day]
  - QTD Change
  - QTD Growth
- YTD | [Quarter] | [Week] | [Day]

- Prior YTD | [Quarter] | [Week] | [Day]
- YTD Grouped
  - Prior YTD | [Quarter] | [Week] | [Day]
  - YTD | [Quarter] | [Week] | [Day]
  - YTD Change
  - YTD Growth

Weekday/Start The weekday when the job started running of which you want sampled, based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour/Start The type of business hour when the job started running, based on the date and time according to the LSF master host.

- Core hours
- Non-core hours

Note: You can define the core hours in the console, under **Data Mappings > Core Hours**.

License Server ◆ [License Server]

Project ◆ [User-defined Project]

User ◆ All Users| [User]

Queue ◆ [Priority] | [Level]

Host Type Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Type] | [Host]

Host Model Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Model] | [Host]

Host Type (Submit) Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Submission Host Type] | [Host]

Host Model (Submit) Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Submission Host Model] | [Host]

Hosts (Submit) Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

	◆ [Submission Host Type] | [Host]
Job Type	◆ [User-defined job]
Job Group	Name of the job group as defined in Platform LSF.
	◆ [Job Group]
# of CPUs	The number of CPUs used to run the job. This is not necessarily the same as the number of CPUs on the host.
	◆ 0.0 (The job was never executed.)
	◆ 1.0
	◆ 4.0
Job Exit Status	◆ DONE
	◆ EXIT
Job Exit Code	◆ [Job Exit Code]
Pending Time Rank	◆ 0 sec to < 5 sec
	◆ 5 sec to < 30 sec
	◆ 30 sec to < 5 min
	◆ 5 min to < 60 min
	◆ 60 min to < 1 day
	◆ One day or more
Vendor	◆ All vendors | [License vendor] | [License feature]
Event	◆ DENIED
	◆ USED
Event Reason	◆ Licensed number of users already reached
	◆ Normal
Job ID	The Job ID as assigned in Platform LSF.
	◆ [Job ID]

Note: To use this dimension, you must enable the "Include Job ID in Workload Cubes" feature in the Platform Analytics Console (click **Addons** in the navigation tree).

Measures

The Hourly Workload License cube includes the following measures:

# of jobs

Total number of jobs.

Job run time (min)

Total run time of all jobs, in minutes.

Job pend time (min)

Total pend time of all jobs, in minutes.

Job turnaround time (min)

Total time elapsed between job submission and job finish.

CPU time (sec)

Total CPU time used by all jobs, in seconds.

Used mem (MB)

Total memory used by all jobs, in megabytes.

Used swap (MB)

Total swap space used by all jobs, in megabytes.

CPU time/Job run time (%)

Total real CPU time used by a job, divided by the total time elapsed between job start and job finish.

Avg run time/job (min)

Average run time of a job, in minutes.

Avg pend time/job (min)

Average pend time of a job, in minutes.

Avg turnaround time/job (min)

Average time elapsed between job submission and job finish for all specified jobs.

Job run time/Turnaround time (%)

Total time elapsed between job start and job finish, divided by the total time elapsed between job submission and job finish.

Avg CPU time/job (min)

Average CPU time used by a job, in minutes.

Avg used mem/job (MB)

Average memory used by a job, in megabytes.

Avg used swap/job (MB)

Average swap space used by a job, in megabytes.

# of checkouts

Number of license checkouts.

Max # of checkouts

Top number of license checkouts at one time for each job; measures the frequency of checkouts for a single application.

Total feature time (min)

Total license feature time per job, in minutes; combines all consumed license time.

Avg feature time (min)

Average license feature time per job, in minutes; averages the consumed license time.

of events

Number of events (for example, denial, support, check in, check out).

Cognos Cube Views

The Hourly Workload License cube does not include any Cognos cube views.

53

The Hourly HPC Workload Cube

The Hourly HPC Workload cube lets you explore how much hardware and time different users and projects consume, and discover whether cluster usage actually maps to your business goals and strategies. It contains historical workload data about HPC jobs, queues, projects, and hosts. It is based on data collected in the Host Properties and HPC Job Information datamarts. You can explore the workload for all HPC projects, queues, users, and hosts.

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change
: The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth
: The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster
: A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hourly HPC Workload cube includes the following dimensions and child categories:

Job Types by Execution Host
: ◆ [Cluster] | [Host] | [Job Type]

Date
: The date when the job finished running, based on the date and time according to the LSF master host.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
◆ [Yesterday]
◆ WTD | [Day]
◆ MTD | [Week] | [Day]
◆ QTD | [Month] | [Week] | [Day]
◆ YTD | |[Month] | [Week] | [Day]
◆ [Previous Week] | [Day]
◆ [Previous Month] | [Week] | [Day]
◆ [Previous Quarter] | [Month] | [Week] | [Day]
◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Jobs by Job Types
: ◆ [Cluster Name] | [Execution Type] | [Job ID] | [Job Array] | [HPC Array]

Weekday
: The weekday when the job was finished.

◆ Sunday
◆ Monday

- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- ◆ Core hours
- ◆ Non-core hours

Shift Based on core hours as defined in the console, under **Data Mappings > Shifts**.

- ◆ Shift 1 | [Hour]
- ◆ Shift 2 | [Hour]
- ◆ Shift 3 | [Hour]
- ◆ By Hour | [Hour]

Date (Submit) The date when the job was submitted, based on the date and time according to the LSF master host.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ Current Day | [Day]
- ◆ Yesterday | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ Prior QTD | [Month] | [Week] | [Day]
- ◆ QTD Grouped
  - ❖ Prior QTD | [Month] | [Week] | [Day]
  - ❖ QTD | [Month] | [Week] | [Day]
  - ❖ QTD Change
  - ❖ QTD Growth
- ◆ YTD | [Quarter] | [Month] | [Week] | [Day]
- ◆ Prior YTD | [Quarter] | [Month] | [Week] | [Day]
- ◆ YTD Grouped
  - ❖ Prior YTD | [Quarter] | [Month] | [Week] | [Day]
  - ❖ YTD | [Quarter] | [Month] | [Week] | [Day]
  - ❖ YTD Change
  - ❖ YTD Growth
- ◆ MTD | [Week] | [Day]
- ◆ Prior MTD | [Week] | [Day]
- ◆ MTD Grouped
  - ❖ Prior MTD | [Week] | [Day]
  - ❖ MTD | [Week] | [Day]
  - ❖ MTD Change
  - ❖ MTD Growth
- ◆ WTD | [Day]
- ◆ Prior WTD | [Day]

- ◆ WTD Grouped
  - ❖ Prior WTD | [Day]
  - ❖ WTD | [Day]
  - ❖ WTD Change
  - ❖ WTD Growth

Weekday (Submit)   The weekday when the job was submitted, based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour (Submit)   The type of business hour when the job was submitted, based on the date and time according to the LSF master host.

- ◆ Core hours
- ◆ Non-core hours

> **Note:** Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

Date (Start)   The date when the job started running, based on the date and time according to the LSF master host.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ Current Day | [Day]
- ◆ Yesterday | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ Prior QTD | [Month] | [Week] | [Day]
- ◆ QTD Grouped
  - ❖ Prior QTD | [Month] | [Week] | [Day]
  - ❖ QTD | [Month] | [Week] | [Day]
  - ❖ QTD Change
  - ❖ QTD Growth
- ◆ YTD | [Quarter] | [Month] | [Week] | [Day]
- ◆ YTD Grouped
  - ❖ YTD | [Quarter] | [Month] | [Week] | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ Prior MTD | [Week] | [Day]
- ◆ MTD Grouped
  - ❖ Prior MTD | [Week] | [Day]
  - ❖ MTD | [Week] | [Day]
  - ❖ MTD Change

- ❖ MTD Growth
- ◆ WTD | [Day]
- ◆ Prior WTD | [Day]
- ◆ WTD Grouped
  - ❖ Prior WTD | [Day]
  - ❖ WTD | [Day]
  - ❖ WTD Change
  - ❖ WTD Growth

Weekday (Start) The weekday when the job started running, based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Business Hour (Start) The type of business hour when the job started running, based on the date and time according to the LSF master host.

- ◆ Core hours
- ◆ Non-core hours

Note: Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

Project
- ◆ [User-defined Project]

User
- ◆ [All Users] | [User]

Queue
- ◆ [Priority] | [Level]

User Group Name of user group as defined in Platform LSF.

- ◆ [User Group]

Note: If a user belongs to multiple user groups, then the job submitted by that user is not associated with a group unless "-G"' is used to indicate the group when submitting the job.

Job Type
- ◆ [Priority] | [Level]

Exit Status
- ◆ Standard
  - ❖ DONE
  - ❖ EXIT

Exit Code
- ◆ Standard
  - ❖ [Job Exit Code]

Pending Time Rank
- ◆ 0 sec to < 5 sec
- ◆ 5 sec to < 30 sec

	◆ 30 sec to < 5 min
	◆ 5 min to < 60 min
	◆ 60 min to < 1 day
	◆ One day or more
# of CPUs per Host	The number of physical processors in the execution host.
	◆ 1
	◆ 2
	◆ 4
# of Cores per CPU	The number of cores in each physical processor in the host.
	◆ 1
	◆ 2
	◆ 4
# of Threads per Core	The maximum number of threads in each core in the host.
	◆ 1
	◆ 2
Application Tag	◆ [Application Tag]

Measures

The Hourly Workload cube includes the following measures:

of jobs

Total number of jobs.

Job run time (min)

Total run time of all jobs, in minutes.

Job pend time (min)

Total pend time of all jobs, in minutes.

Job turnaround time (min)

Total time elapsed between job submission and job finish.

CPU time (sec)

Total CPU time used by all jobs, in seconds.

Used mem (MB)

Total memory used by all jobs, in megabytes.

Used swap (MB)

Total swap space used by all jobs, in megabytes.

CPU time/Job run time (%)

Total real CPU time used by a job, divided by the total time elapsed between job start and job finish.

Number of Bits

Total number of job bits.

Number of Execution Processes

Total number of execution processes for the job.

Cognos Cube Views

The Hourly Workload cube includes the following Cognos cube views:

- Consumed resources by job exist status and cluster
- Jobs ranked by pending time
- Throughput efficiency by cluster and date
- Types of hosts completing jobs the fastest
- Projects consuming the most CPU time
- Projects running the most jobs
- Projects waiting the longest for cluster resources
- Projects wasting the most CPU time
- Projects with the most CPU-intensive jobs
- Throughput efficiency by cluster and project
- Jobs in queue ranked by pending time
- Number of done and exited jobs by queue
- Throughput efficiency by queue
- Users consuming the most CPU time
- Users waiting the longest to execute jobs
- Users with the most jobs

54

The License Utilization Prediction Cube

The License Utilization Prediction cube contains historical and predictive data about license utilization from a license server perspective. It is based on the License Utilization Prediction datamart. It does not contain information about individual users.

Contents

Dimensions

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The License Utilization Prediction cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Weekday   The predicted weekday.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

License server
- [License Server]

Measures

The License Utilization Prediction cube includes the following measures:

Average utilization (%)

The average license utilization on each license server over a period of time, expressed as a percentage.

Peak utilization (%)

The peak license utilization on each license server over a period of time, expressed as a percentage.

4th week forecasted avg utilization (%)

Average license utilization predicted for this host in the 4th week, expressed as a percentage.

4th week forecasted peak utilization (%)

Maximum license utilization predicted for this host in 4 weeks, expressed as a percentage.

8th week forecasted avg utilization (%)

Average license utilization predicted for this host in the 8th week, expressed as a percentage.

8th week forecasted peak utilization (%)

Maximum license utilization predicted for this host in 8 weeks, expressed as a percentage.

12th week forecasted avg utilization (%)

Average license utilization predicted for this host in the 12th week, expressed as a percentage.

12th week forecasted peak utilization (%)

Maximum license utilization predicted for this host in 12 weeks, expressed as a percentage.

Cognos Cube Views

The License Utilization Prediction cube includes the following Cognos cube view report:

- Peak utilization trend by feature and server over the next 12 weeks

55

The Job Throughput Prediction Cube

The Job Throughput Prediction cube contains historical and predictive job throughput data. It is based on data collected from the `lsb.acct` file, collected in the Job Throughput Prediction datamart. The dynamic historical data is collected hourly and aggregated daily. The cube is also based on the Job Throughput Prediction datamart, which contains predictive data.

Contents

Dimensions

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Job Throughput Prediction cube includes the following dimensions and child categories:

Weekday   The weekday when the job is predicted to finish running.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Cluster
- [Cluster] | [Host]

Queue
- [Priority] | [Level]

Measures

Data collection

Data about job throughput is collected in "real time", so that it does not include data about running or pending jobs. Therefore, jobs that were submitted before midnight but are still running or pending at the last midnight before the cube build are not included in the count of submitted jobs. Consequently, the data in the cubes data changes frequently, especially when there are lots of long jobs, which run for more than one day.

Consider the example of a single cluster that is monitored by Analytics with the following average hourly statistics:

- 300 jobs are submitted
- 200 jobs finish running with done status
- 50 jobs finish running with exit status
- 20 jobs never run
- 30 jobs are still running when Analytics analyzes the data

The average number of submitted jobs in the resulting cube will increase from the time Analytics originally analyzes the data to once the 30 running jobs finish running.

Calculation of measures

While all displayed measure values are hourly averages, they are actually an aggregate of 24 hourly averages, divided by 24. You cannot drill into the data to see this data for a particular hour in the day.

Measures

The Job Throughput Prediction cube includes the following measures:

Average # of submitted jobs

Total number of submitted jobs.

Average # of exited jobs

Total number of exited jobs.

Average # of done jobs

Total number of done jobs.

4th week forecasted # of submitted jobs

Total number of submitted jobs.

4th week forecasted# of exited jobs

Total number of exited jobs.

4th week forecasted# of done jobs

Total number of done jobs.

8th week forecasted # of submitted jobs

Total number of submitted jobs.

8th week forecasted # of exited jobs

Total number of exited jobs.

8th week forecasted # of done jobs

Total number of done jobs.

12th week forecasted # of submitted jobs

Total number of submitted jobs.

12th week forecasted # of exited jobs

Total number of exited jobs.

12th week forecasted # of done jobs

Total number of done jobs.

Cognos Cube Views

The Job Throughput Prediction cube includes the following Cognos cube views:

- Submitted job trend per cluster over the next 12 weeks
- Submitted job trend per queue over the next 12 weeks

56

The Jobs vs. Slot Utilization Cube

The Jobs vs. Slot Utilization cube lets you determine the following:

◆ The number of slots were available on a cluster or a group of clusters.

◆ The number of slots were being used and the number of jobs running.

◆ The number of pending slots and jobs in the system.

This cube allows you to make conclusions about the demand on the system and how they are utilized. For example,

◆ If the number of available slots is consistent, the cluster is stable. Conversely, if the number of available slots changes frequently, that means there are hosts going offline and coming online, which indicates an unstable cluster..

◆ If the number of used slots is close to the number of available slots, this indicates that your cluster has good utilization.

◆ If the number of used slots is much smaller than the number of available slots, this indicates your cluster has spare capacity. If there are lots of pending jobs and slots, this indicates a potential problem in your cluster. and you should examine why the jobs or slots are pending.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Jobs vs. Slot Utilization cube includes the following dimensions and child categories:

Category
- SUSP
- WAIT
- RUN
- PEND
- CAPACITY

Cluster   Cluster | [Host]

Date   The sampling date and time.

Note: The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Host Type Host type information of the job execution host is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Type] | [Host]

Host Model Host model information of the job execution host is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Model] | [Host]

Slot Qty The number of execution processors of a slot that is using the CPU.

Project
- [User-defined Project]

Queue
- [Queue]

User
- [User]

Weekday Based on the date and time according to the LSF master host.

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Measures

The Jobs vs. Slot Utilization cube includes the following measures:

# of Slots   The average number of slots over time.

# of Jobs   The average number of jobs over time.

Cognos Cube Views

The Jobs vs. Slot Utilization cube does not include any Cognos cube views.

57

The Cluster Capacity Cube

The Cluster Capacity cube provides a summary of how one or more clusters are being used by allowing you to view the size, availability, usage, and capacity of the cluster (or clusters) in question.

This cube allows you to determine the following:

◆ The overall state of the slots in each cluster: how many slots are available, how many slots are closed (due to administration), and how many slots are down (due to hosts being unavailable).

◆ The number of idle slots in each cluster.

◆ Which workload is running on the slots in each cluster, organized by project, queue, user, host type, or host model.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Cluster Capacity cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Host Type   Host type information of the job execution host is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Type] | [Host]

Host Model   Host model information of the job execution host is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.

- [Host Model] | [Host]

Project	◆ [User-defined Project]
Queue	◆ [Queue]
User	◆ [User]
Job CPU Qty	The number of execution processors of a job that is using the CPU.
Weekday	Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday
- ◆ Thursday
- ◆ Friday
- ◆ Saturday

Measures

The Cluster Capacity cube includes the following measures:

Avg Slots   The average number of slots used over time.

Sum Slots   The total number of slots used over time.

Cognos Cube Views

The Cluster Capacity cube does not include any Cognos cube views.

58

The Health Check Cube

The Health Check cube reports on how old the data is for each specified cube in each cluster that you are monitoring. This allows you to see when the data is lagging behind.

The data may lag behind because the loader controller or the Analytics node is down, the corresponding data loaders are slow, or the data aggregation has an extended latency. Using this cube, you can check data latency, which data flow is experiencing problems, and by how many days a problematic data flow is behind.

By default, the Health Check Cube monitors the following cubes:

- ◆ Hourly/Daily Hardware cube
- ◆ Hourly/Daily Workload cube
- ◆ Hourly/Daily Job Throughput cube

Contents

Dimensions

Abbreviations

WTD	Week to Date
MTD	Month to Date
QTD	Quarter to Date
YTD	Year to Date

Definitions

Change	The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.
Growth	The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.
Cluster	A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Health Check cube includes the following dimensions and child categories:

Cluster	Cluster	[Host]
Cube Name	Cube Name

Measures

The Health Check cube includes the following measures:

Latency The number of days delayed. The Health Check cube derives this value by subtracting the latest time of data in the source table from the cube build time.

Cognos Cube Views

The Health Check cube does not include any Cognos cube views.

59

The Hardware Histogram Cube

The Hardware Histogram cube lets you explore the load on your cluster, and to ensure that it is configured correctly and fully utilized. It contains hourly historical hardware data, and it is based on the Hardware Histogram datamart. Metrics are aggregated for groups or types of hosts using a weighted average. You can view data by year, quarter, month, week, day, or hour. You can explore detailed hardware usage for all hosts, host models, host types, and rank types.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Cluster   A group of computers (hosts) running LSF that work together as a single unit, combining computing power and sharing workload and resources.

Dimensions

The Hardware Histogram cube includes the following dimensions and child categories:

Cluster   Cluster | [Host]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- ◆ [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- ◆ [Yesterday]
- ◆ WTD | [Day]
- ◆ MTD | [Week] | [Day]
- ◆ QTD | [Month] | [Week] | [Day]
- ◆ YTD | |[Month] | [Week] | [Day]
- ◆ [Previous Week] | [Day]
- ◆ [Previous Month] | [Week] | [Day]
- ◆ [Previous Quarter] | [Month] | [Week] | [Day]
- ◆ [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.

- ◆ Sunday
- ◆ Monday
- ◆ Tuesday
- ◆ Wednesday

	◆ Thursday
	◆ Friday
	◆ Saturday
Business Hour	Based on core hours as defined in the console, under **Data Mappings > Core Hours**.
	◆ Core hours
	◆ Non-core hours
Shift	Based on shift hours as defined in the console, under **Data Mappings > Shifts**.
	◆ Shift 1 | [Hour]
	◆ Shift 2 | [Hour]
	◆ Shift 3 | [Hour]
	◆ By Hour | [Hour]
Host Model	Host model information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Model] | [Host]
Host Type	Host type information is based on the LSF master LIM (Load Information Manager), which is collected and updated hourly.
	◆ [Host Type] | [Host]
Rank Type	The rank type for this cube.
	◆ Mem util (%) rank
	The memory utilization rank.
	❖ 0~10%
	❖ 10~20%
	❖ 30~40%
	❖ 40~50%
	❖ 50~60%
	❖ 60~70%
	❖ 70~80%
	❖ 80~90%
	❖ 90~100%
	◆ CPU util (%) rank
	The CPU utilization rank.
	❖ 0~10%
	❖ 10~20%
	❖ 20~30%
	❖ 30~40%
	❖ 40~50%
	❖ 50~60%
	❖ 60~70%
	❖ 70~80%
	❖ 80~90%

❖ 90~100%

Measures

The Hardware Histogram cube includes the following measures:

Occurrence   The number of occurences for the given rank.

Duration   The length of time (in minutes) for the given rank. This is calculated by the following formula for the given rank: **Occurrence** $\times$ **interval_period**.

Cognos Cube Views

The Hardware Histogram cube does not include any Cognos cube views.

60

The License Usage by Server Histogram Cube

The License Usage by Server Histogram cube lets you explore the license usage from the license server perspective. It contains historical license utilization data, and it is based on the License Usage by Server Histogram datamart. It does not contain information about individual users.

Contents

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The License Usage by Server Histogram cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- Core hours
- Non-core hours

License server
- [License Server]

Rank Type The rank type for this cube.

- Utilization (%) rank
  The license utilization rank.
    - 0~10%
    - 10~20%
    - 20~30%
    - 30~40%
    - 40~50%
    - 50~60%
    - 60~70%
    - 70~80%
    - 80~90%
    - 90~100%
- Used number rank
  The number of licenses used.
    - 0~10 licenses
    - 10~20 licenses
    - 20~80 licenses
    - 80~100 licenses
    - 100 licenses or more

Measures

The License Usage by Server Histogram cube includes the following measures:

Occurrence   The number of occurences for the given rank.

Duration   The length of time (in minutes) for the given rank. This is calculated by the following formula for the given rank: **Occurrence** $\times$ **interval_period**.

Cognos Cube Views

The Hourly License Utilization cube does not include any Cognos cube views.

61

The License Usage by Feature Histogram Cube

The License Usage by Feature Histogram cube lets you explore license usage information by feature. It is based on the License Usage by Feature Histogram datamart.

Contents
- "Dimensions" on page 376
- "Measures" on page 378
- "Cognos Cube Views" on page 379

Dimensions

Abbreviations

WTD   Week to Date

MTD   Month to Date

QTD   Quarter to Date

YTD   Year to Date

Definitions

Change   The change from the previous specified time category to the current time category. For example, QTD Change is the change from the previous QTD to the current QTD.

Growth   The change from the previous specified time category to the current time category, expressed as a percentage. For example, QTD Growth is the change from the previous QTD to the current QTD, expressed as a percentage.

Dimensions

The License Usage by Feature Histogram cube includes the following dimensions and child categories:

License vendor
- All vendors | [License vendor] | [License feature]

Date   The sampling date and time.

> **Note:** The week starts on Sunday and ends on Saturday. The Date dimension is the internal timestamp.

- [Current Year] | [Quarter] | [Month] | [Week] | [Day]
- [Yesterday]
- WTD | [Day]
- MTD | [Week] | [Day]
- QTD | [Month] | [Week] | [Day]
- YTD | |[Month] | [Week] | [Day]
- [Previous Week] | [Day]
- [Previous Month] | [Week] | [Day]
- [Previous Quarter] | [Month] | [Week] | [Day]
- [Previous Year] | [Quarter] | [Month] | [Week] | [Day]

Weekday   Based on the date and time according to the LSF master host.
- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Business Hour   Based on core hours as defined in the console, under **Data Mappings > Core Hours**.

- Core hours
- Non-core hours

Rank Type   The rank type for this cube.

- Utilization (%) rank
  The license utilization rank.

  - 0~10%
  - 10~20%
  - 20~30%
  - 30~40%
  - 40~50%
  - 50~60%
  - 60~70%
  - 70~80%
  - 80~90%
  - 90~100%

- Used number rank
  The number of licenses used.

  - 0~10 licenses
  - 10~20 licenses
  - 20~80 licenses
  - 80~100 licenses
  - 100 licenses or more

Measures

The License Usage by Feature Histogram cube includes the following measures:

Occurrence The number of occurences for the given rank.

Duration The length of time (in minutes) for the given rank. This is calculated by the following formula for the given rank: **Occurrence** $\times$ **interval_period**.

Cognos Cube Views

The License Usage by Feature Histogram cube does not include any Cognos cube views.

Index