

ORACLE[®]

PRIMAVERA

**P6 Analytics and P6 Reporting Database Planning and Sizing Guide
Release 3.2**

December 2013

Contents

Overview of Planning Your Implementation	7
About Oracle Primavera P6 Enterprise Project Portfolio Management.....	7
About P6 Analytics.....	8
About P6 Reporting Database	9
Working with Oracle Primavera P6 EPPM, P6 Analytics, and P6 Reporting Database.....	11
Critical Performance Factors	12
Pulling Data between Servers.....	12
Merging Updates into Target Database	13
PL/SQL-based Transformations.....	13
Planning Process	13
Requirements Phase	13
Using Oracle Database Enterprise Edition (without partitioning) for the Star Schema .	14
Key Questions to Ask about Operational Reporting.....	14
Key Questions to Ask about P6 Analytics/Business Intelligence.....	15
Planning Phase.....	15
Full ETL Process Only.....	15
Monitoring Usage	16
Basic Monitoring.....	16
Physical Sizing	17
Overview of Physical Space Usage	17
Table Partition for Multiple Data Source and History Data	17
Archiving and Restoring History	18
Archiving History.....	18
Restoring Archived History	19
Physical Components	20
ETL Process Server	20
Calculating Spread Sizes	20
Queries for Spread Sizing.....	20
Estimating the Size of the Star Database.....	22
Estimating the Size of the ODS Database.....	22
Oracle Star Database.....	23
Physical Hardware.....	24
Monitoring Dimension and History Level Impacts on the Database	24
Consideration for ETL Performance	24
Network.....	25
Memory.....	25
Extract and Load.....	26
Summary of Physical Hardware Sizing.....	26
Size of P6 EPPM Database.....	26
Amount of Change.....	26
Complexity of Project Data	26

Expectations and Operations	26
Small Database Deployment	27
Medium Database Deployment	27
Large Database Deployment	27
Database Servers	28
Planning Revisited	28
Timing Estimation	28
Factors Affecting ETL Runtime	30
Planning for P6 EPPM's Impact on Star	32
Choosing Between Standard Edition and Enterprise Edition	32
P6 Extended Schema's Affect on Star	33
Date Range	33
Enable Publication	33
History Settings on Project Level	34
Slowly Changing Dimensions	34
Multiple Star Environments	35
Creating Multiple Data Sources for a Single Star Database	35
Setting up Multiple Data Sources	36
Filtering by Data Source in Analytics	36
Single Star, Single Data Source, Multiple ETLs	37
Setting up Multiple ETLs	37
Run the ETLs	37
Appendix: Checklists	38
Checklist for P6 Analytics and the Star Database	38
Pre-installation for an Oracle Platform	38
Pre-installation for a Microsoft SQL Server Platform	39
Pre-installation for Configuring the ETL Server	39
Pre-installation for the Star Database	40
Installation: Installer And Configuration	42
Installation: Running The Processes	44
Post-installation: After The Processes Complete	44
Database Scripts For Validations	44
Checklist for ODS Database	46
Pre-installation for an Oracle Platform	46
Pre-installation for a Microsoft SQL Server Platform	46
Pre-installation for Configuring the ETL Server	47
Pre-installation for the ODS Database	48
Installation: Installer And Configuration	49
Installation: Running The Processes	50
Post-installation: After The Processes Complete	51
Database Scripts For Validations	51
For More Information	53
Where to Get Documentation	53
Where to Get Training	55

Where to Get Support.....55

Legal Notices.....57

Overview of Planning Your Implementation

P6 Reporting Database and P6 Analytics are data warehousing products. Implementing them involves databases, servers, and controlling an ETL process. Oracle Primavera provides the code to perform the ETL process. The P6 Analytics module contains a pre-built application including dashboards, reports, and underlying metadata to provide end users with easier access to information.

When planning each deployment, you need to consider:

- ▶ What's required for physical storage and CPU processing.
- ▶ Which components of P6 Reporting Database and P6 Analytics will be implemented.
- ▶ The differences in data volumes, frequency of data changes, and business requirements.

To successfully implement these products:

- 1) Review the critical performance factors.
- 2) Outline a methodology for planning an installation.
- 3) Determine the physical storage requirements of the data warehouse.
- 4) Address the server performance requirements of the ETL process.

In This Section

About Oracle Primavera P6 Enterprise Project Portfolio Management.....	7
About P6 Analytics	8
About P6 Reporting Database	9
Working with Oracle Primavera P6 EPPM, P6 Analytics, and P6 Reporting Database	11

About Oracle Primavera P6 Enterprise Project Portfolio Management

Oracle Primavera P6 Enterprise Project Portfolio Management (P6 EPPM) is a robust and easy-to-use integrated solution for globally prioritizing, planning, managing, and executing projects, programs, and portfolios. It optimizes role-specific functionality to satisfy each team member's needs, responsibilities, and skills. It provides a single solution for managing projects of any size, adapts to various levels of complexity within a project, and intelligently scales to meet the needs of various roles, functions, or skill levels in your organization and on your project team.

Thousands of companies rely on P6 EPPM to:

- ▶ Plan, schedule, and manage the right strategic mix of projects
- ▶ Make informed project, cost, and resource management decisions
- ▶ Enhance team collaboration and workflow productivity
- ▶ Maximize each project contributor's skills and responsibilities
- ▶ Track progress and optimize capacity for maximum profitability
- ▶ Deliver programs and projects on time and within budget

- ▶ Share data with human capital, financial management, and enterprise resource planning systems

About P6 Analytics

P6 Analytics provides an in-depth and comprehensive method for analyzing and evaluating project performance, project history, resource assignments, and utilization.

Built upon the Oracle Business Intelligence (OBI) suite, it delivers a catalog of analyses that provide an interactive way of viewing, analyzing, and evaluating P6 EPPM data. In addition, it provides a Repository Definition (.rpd) file that contains the data mappings between the physical data and the presentation layer of OBI.

The dashboards provide detailed insight into your P6 EPPM data through analytical charts, tables, maps, and graphics. Dashboards allow you to navigate to other analyses to provide precise root-cause analysis. OBI allows you to configure individual analyses with the P6 EPPM Action Link, enabling you to navigate directly to your P6 site for true "Insight to Action" capabilities. You can save analyses created with OBI Answers in the OBI Presentation Catalog and integrate the analyses into any OBI home page or dashboard. You can enhance results through options such as charting, result layout, calculation, and drill-down features.

Use P6 Analytics to:

- ▶ Perform root-cause analysis and employ management-by-exception.
- ▶ Gather critical insights into current and historical performance of all projects, programs, and portfolios.
- ▶ Make better decisions to eliminate project failure.
- ▶ Quickly visualize critical project performance in early stages.
- ▶ Predict and identify cost trends early in the project life cycle to rescue troubled projects.
- ▶ Gain visibility into resource performance through s-curves. With interactive dashboards, you can drill down to examine the root-cause of a problem.
- ▶ Show staffing needs by portfolio with early warning indicators for upcoming under-staffed project work.
- ▶ Use geospatial visualization to view project, activity, and resource data metrics by geographic location with full drill-down capabilities.

Performance Data

P6 Analytics provides an .rpd file to use with the OBI suite. The .rpd file contains:

- ▶ A physical representation of the Star schema.
- ▶ A business layer to perform customized calculations.
- ▶ A presentation layer that groups all the calculated business layer fields into logical subject areas.

The .rpd delivers an extensive amount of Earned Value, Costs, Units, Percent Completes, and other key performance indicators. It enables data to be sliced by dimensions such as time, EPSs, portfolios, projects, activities, and resources.

P6 Analytics delivers a sample dataset, consisting of Star data, where the dashboards and analyses in the catalog were built. You can use this sample data to view the power of dashboard and analyses delivered in the catalog, and see how you can integrate the catalog with your data.

About P6 Reporting Database

P6 Reporting Database works with the P6 EPPM database to provide a robust and powerful reporting solution. P6 Reporting Database consists of the Star and ODS databases.

The Star Database

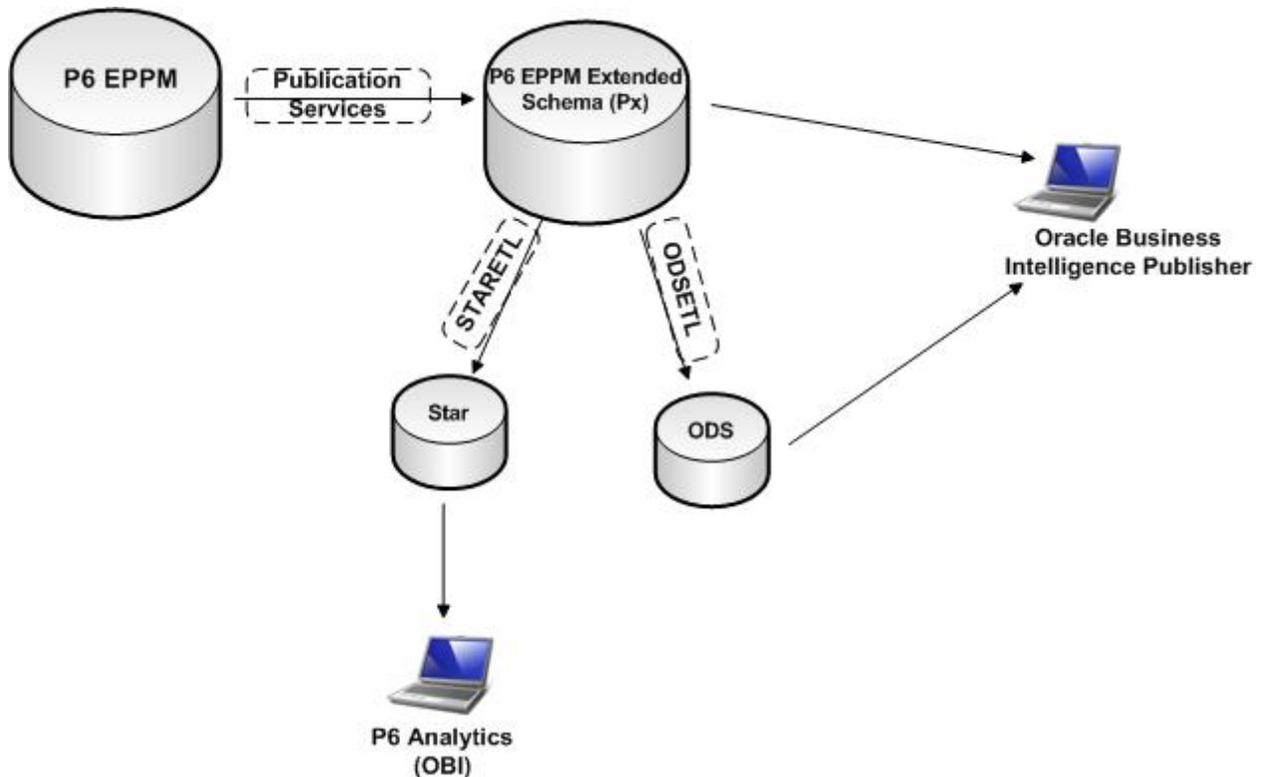
The Star database enables an organization to perform advanced business analysis on project and portfolio data. It supplies a dimensional schema that organizes P6 EPPM hierarchical relationships. The Star schema allows organizations to store data for History Levels (Project, WBS, Activity, or Assignment) for a specific History Interval (Year, Quarter, Financial Period, Month, Week, or Day). This data allows for tracking trends for advanced business intelligence. The Star database captures data for activities, resource assignments and utilization, and user defined fields. P6 Analytics requires the Star database.

The Operational Data Store (ODS)

The ODS portion of P6 Reporting Database is an optional relational database that supplies day-to-day, easy to understand operational views of the P6 EPPM database data. You can also use the P6 Extended Schema to provide this information.

Working with Oracle Primavera P6 EPPM, P6 Analytics, and P6 Reporting Database

Depending on your organization's specific deployment, P6 EPPM, P6 Reporting Database, and P6 Analytics work together as depicted below.



Applications:

P6 EPPM: P6 EPPM is the primary interface for administering and managing projects. This is where you will gather the data on which P6 Analytics will report. P6 EPPM provides an extensive array of features designed to optimize all phases of Project Management, Resource Management, and Portfolio Management. It includes full support for activities, work breakdown structures, planning and scheduling, costs, resource administration and assignment, codes, user defined fields, roles, teams, and capacity planning.

Oracle Business Intelligence (OBI): OBI integrates data from multiple enterprise sources, and transforms it into key insights that enable strategic decision-making, drive continual business process improvements, and promote alignment across the enterprise. OBI can also generate proactive alerts based on integrated reports data mined from P6 Reporting Database.

Oracle BI Publisher: Oracle Business Intelligence Publisher is an enterprise reporting solution allowing you to design, manage, and deliver highly formatted documents. Because it is built on open standards, your IT developers can create data models against practically any data source and use BI Publisher APIs to build custom applications leveraging existing data sources and infrastructure. BI Publisher users can design report layouts using familiar desktop tools, reducing the time and cost needed to develop and maintain reports. Extremely efficient and highly scalable, BI Publisher can generate documents with minimal impact to transactional systems.

P6 Analytics: This integrated dynamic reporting tool, uses the data gathered in P6 and the OBI metadata layer, to facilitate the creation of ad hoc reports and interactive custom dashboards reflecting trends and metrics for activities, portfolios, resource assignments, utilization, and project history.

Technology:

The P6 EPPM Database: The main database for P6 EPPM data. The EPPM database includes the P6 Extended Schema (Px) optimized for reporting.

The Reporting Database: The P6 Reporting Database consists of the Star database and the Operational Data Store (ODS) database. Each of these pulls data from the P6 EPPM Extended Schema, which can then be used to create reports. The Star database accumulates project data over time, providing baselines for tracking trends and for advanced business intelligence; it also provides key metrics on current data, which allows you to drill through root cause analysis. Users can view the data by using OBI. The ODS database gives users the option to access denormalized, relational views of P6 EPPM data by using tools such as Oracle Business Intelligence Publisher.

Critical Performance Factors

The ETL process has four areas that will affect performance:

- 1) Pulling data between servers.
- 2) Managing the updates of the component databases.
- 3) Performing PL/SQL and direct SQL transformation operations on the database server.
- 4) Project management data dimension and fact table generations.

See the Tested Configurations document to determine what product versions P6 Analytics and P6 Reporting Database

Pulling Data between Servers

Data movement revolves around the E (Extract), T (Transform), and L (Load) in ETL. In a typical implementation, you will deploy the P6 EPPM Database and data warehouse (ODS/STAR) on separate physical servers. So, you must maximize the bandwidth and minimize latency between servers. Ideally, the servers reside in the same data center with gigabit ethernet connecting the servers. (Oracle Primavera performance tests are performed with servers in a central data center with gigabit connections.) You should verify throughput of server communication by performing basic file copy or FTP operations between the servers.

The data movement processes is based on the standard SQL, either with INSERT-SELECT or CREATE TABLE AS (CTAS) syntax using an Oracle database link. While the Oracle RDBMS is efficient at moving data through the link, the overall performance depends on the physical network connection.

Merging Updates into Target Database

P6 Analytics and P6 Reporting Database leverage PL/SQL Bulk operations to process smaller batches while performing interim commits. This results in linear scaling of update operations (for example, if it takes five minutes to update one million rows, it will take ten minutes to update two million rows). Rows per second performance remains constant here, regardless of amount of change or data size.

Many of the processes, including the PL/SQL Bulk operation processes, are run in parallel threads. The use of parallel PL/SQL processes not only increases the potential scalability of this process, but also the demand for CPU powers and network bandwidth on the server.

PL/SQL-based Transformations

PL/SQL does a portion of the transformation process. The largest portion of this is referred to as direct SQL. SQL update statements run directly against tables to perform simple transformation that does not require business logic. These processes are generally executed in parallel, and are very CPU intensive, on the database server. With the large chunk of data processed, the disk IO performance is limiting. A high performing SAN system improves the overall system response time.

Planning Process

The planning process includes a Requirements Phase, where you evaluate system requirements, and Planning Phase, where you evaluate impacts on your system.

Requirements Phase

The first phase is to understand what the users of the system want to get from the solution. This includes:

- ▶ The types of reports they want
- ▶ The level of detail needed in reports
- ▶ The time frame they need reports delivered
- ▶ How current the data should be

P6 Analytics and P6 Reporting Database provides a wide variety of results. However, you may not need them all for a given installation. There are two categories of reporting solutions:

- ▶ Operational
- ▶ Analytics/Business Intelligence

Using Oracle Database Enterprise Edition (without partitioning) for the Star Schema

If you are connecting multiple P6 data sources, you must use Oracle Database Enterprise Edition with Partitioning (you cannot use Standard Edition).

If your Star schema is installed using Oracle Standard or Enterprise Edition without partitioning, the P6 EPPM database should be a small-sized database as defined in the **Planning for P6 EPPM's Impact on Star** (on page 32) section of this document. Partitioning helps keep performance consistent over time for each ETL run. This is particularly important when enabling activity or WBS-level history. If you choose not to use partitioning, Oracle recommends that you **not** use activity and WBS-level history gathering. If the P6 EPPM database is larger than the defined criteria of a small-sized database, Oracle recommends you use Oracle Database Enterprise Edition with Partitioning.

Without partitioning, the amount of time it takes to run the ETL process will increase over time. Track the row counts in the w_project_history_f, w_wbs_history_f, w_activity_history_f. As the size of these tables increases the need for partitioning will also increase. You can monitor these run times in the staretl.html and staretlprocess.log located in the <installation directory>\log folder. Oracle Database Enterprise Edition with Partitioning is the solution for growth over time within the history tables.

Key Questions to Ask about Operational Reporting

Note: ODS is secondary to the P6 EPPM Extended Schema. The P6 EPPM Extended Schema is the primary resource for operational reporting.

Operational reporting covers the day-to-day reports used by project managers, resource managers, and other tactical personnel. When using operational reporting, you must consider the scheduling and delivery of reports. The combination of the spread data persistence in the P6 EPPM schema and BI Publisher addresses the scheduling, execution, and delivery of the reports. On a day-to-day basis, the reporting load will be fairly consistent. You need to answer the following:

- ▶ **When will reports run?** Users may need to execute on demand. Or they may need reports prior to the start of the work day. These considerations will affect the timeliness of the data. You can run it one or more times per day and limit the duration of each run (which is proportional to the data volume and date range selected by the users).
- ▶ **How will reports be delivered?** You must get the right reports to users at the right time. BI Publisher offers multiple ways to deliver reports from the P6 EPPM Extended Schema. These include email, HTTP, WEBDAV, and direct printing. The logistics of setting up these delivery methods must be considered during the planning process. Moreover, any leading third-party reporting tools can be utilized to generate and deliver the reports.
- ▶ **What will the reporting load be on Oracle ODS Database?** One of the major considerations affecting subsequent decisions will be the load on the reporting server.
- ▶ Queries will execute against the Oracle ODS database to fulfill reporting requests. You must consider that this usage will peak during specific times. Since the exact types of queries are unknown at this point, you should analyze:

- ▶ **How many users are accessing the database at the same time?** This will determine the maximum load on the database server.
- ▶ **Is the reporting on individual projects or across the entire database?** Aggregate queries will put more load on the server than project-specific queries.
- ▶ **Is the reporting done in batch or interactively?** More interactive reporting will increase the demands on both the server CPU and I/O subsystem.

Note: The BI Server component is capable of robust caching of query results, which can mitigate performance concerns. The effectiveness of caching depends on how much users share security. If every user's project access is distinct (including the level of access to cost fields) then the cache will only be effective for each user individually.

Key Questions to Ask about P6 Analytics/Business Intelligence

P6 Analytics reporting is dynamic. The Oracle Star Database schema and OBI integration was designed to allow a very rich environment. This means that the daily load on the data warehouse server and OBI will vary greatly. You should consider:

- ▶ **Who will access P6 Analytics?** A diverse set of users from the CEO to resource/project managers may require access to P6 Analytics. Each may have different requirements and use cases to consider.
- ▶ **What are the default ways of filtering?** By default, user requests for analytic information will include all the data accessible by that user. That may be more time consuming, and may include more information than necessary. Consider ways of filtering data, such as by Project Codes and Portfolios.
- ▶ **What codes are used for reporting?** While the Oracle ODS Database includes data from the P6 EPPM database, the Oracle Star Database schema includes only a subset of codes among activity, resource, project codes, and UDFs. You must determine which codes are critical for analysis.

Planning Phase

This implementation will place high demands on physical storage. Because the calculation process places a unique demand on the ETL, this data warehouse implementation may require higher amounts of CPU/memory. Subsequent sections of this document will provide more detailed information. First, this document will examine the high-level aspects of planning the implementation.

Full ETL Process Only

This release runs only in full ETL mode. Since the data has been pre-processed by the Publication services in the P6 EPPM Extended Schema, direct loads of the data into ODS and Star are more efficient.

Monitoring Usage

You must understand how data compiles in key tables and how this impacts performance of the usage monitoring process. You can gather audit columns on each table (for example, CREATE_DATE) on a daily basis and from the REFRDEL table. You should monitor how much the tables through several weeks and now the periodic peak activity. You must keep track of peak usage times because you will use them to make hardware decisions. Monitor the following tables:

- ▶ PROJECT
- ▶ TASK
- ▶ TASKACTV
- ▶ TASKRSRC
- ▶ TASKMEMO
- ▶ UDFVALUE
- ▶ RSRCHOUR
- ▶ PROJWBS

Basic Monitoring

Use the following query to count changes to the table (insert and update):

Notes:

- The query should run at the end of the day to get all the changes from that day.
 - This query is repeated for all the critical tables.
-

```
SELECT count(*)
FROM <table>
WHERE update_date > trunc(sysdate)
```

Use the following query to reads/selects rows from REFRDEL in a single step:

```
SELECT TABLE_NAME, count(*)
FROM REFRDEL
WHERE delete_date > (sysdate)
GROUP BY TABLE_NAME
ORDER BY TABLE_NAME
```

When necessary, you can use P6 Auditing to get a precise picture of usage.

Physical Sizing

Overview of Physical Space Usage

The physical space requirements of the data warehouse consist of more than just copies of the project management data. Space requirements will vary with the amount of data processed from the P6 EPPM Database. The system uses space for the following types of data:

- ▶ **Core Project Management Data** This includes all the physical fields that exist in the P6 EPPM Database. This is approximately all the data in the P6 EPPM Database.
- ▶ **Logical Fields** The fields become physically stored as part of the P6 EPPM Database in the Extended Schema (Px) tablespace. While this is less than the size of the P6 EPPM Database, it may be as much as 50% of the total space.
- ▶ **Fact/Spread Data** Spread and Fact data total size depend on the number of activities and resource assignments, the average length of activities, and the total window (date range) of the data warehouse. Because of this, it will be treated as a distinct group. It is part of the ODS and is fundamental to the dimensional schema (Star).
- ▶ **Other ETL Tables** There is some database space usage specific to the ETL process. This space is trivial relative to the core P6 EPPM data.
- ▶ **ETL Process Installation** This includes shell scripts, sql files, and JAR files.

Table Partition for Multiple Data Source and History Data

- ▶ **Purpose of Table Partition** Improved performance and ease of data management.
- ▶ **Multiple Data Source** Star supports multiple P6 databases as the data source.
By default, the related tables are value partitioned based on the DATASOURCE_ID. The Partitioned tables are visible by viewing `\scripts\create_star_tables_part.sql` and searching for tables with **PARTITION**.
- ▶ **History Data** For P6 Reporting Database 3.2 you can control history interval and levels settings on the project level in P6. History data can grow quickly. For example, if you select Activity as the History Level, this will automatically set the History Interval for Activity and Resource Assignment history to Daily, which can effect the ETL process' performance.
By default, these tables are interval-partitioned based on the PERIOD_END_DATE value, then sub-partitioned by the value of the DATASOURCE_ID. Oracle recommends you allocate of separate tablespace for each partition. The 3 history tables are:

Project history:

W_PROJECT_HISTORY_F

WBS history:

W_WBS_HISTORY_F

Activity history:

W_ACTIVITY_HISTORY_F

The indexes on these tables are LOCAL indexes per each partition.

- ▶ **Slowly Changing Data for projects with Activity as the History Level** This will automatically set the History Interval for Activity and Resource Assignment history to Daily, which can effect the ETL process' performance. By default, these tables are interval partitioned based on the effective_end_date value, then sub-partitioned by the value of the DATASOURCE_ID. The tables related to this feature end in _HD and _HF.
- ▶ **Number of Partitions** are dynamically determined for the interval partition. For datasource partitioning, 3 are created by default. An example to add more datasource partitions is given below.
 - ▶ **Value Partition Addition:**

```
alter table W_ACTIVITY_D add partition P4 values (4) tablespace
star_hst1;
```
- ▶ **Interval Partition Management** The merging of partitions may be applicable in some environments if the record counts found in each partition are relatively small. In the INSTALL_HOME/script directory a history_partition_management.sql script is provided. This script has examples on how to merge partitions based on an upper thresholds and a function to drop empty partitions if any were created in error. This script serves as a template in partition management.

Archiving and Restoring History

This section takes you through how to archive and restore archived history from the database.

Archiving History

To archive history, you need to take the history partition and move the data out of the database into a file.

To archive history:

Note: The commands below represent the first history partition.

- 1) Drop security on the W_ACTIVITY_HISTORY_F so you can exchange the partition to a table. Run this command to drop security:

```
begin
    dbms_rls.drop_policy('STARUSER', 'W_ACTIVITY_HISTORY_F',
'W_ACTIVITY_HISTORY_F_P_POLICY');
    dbms_rls.drop_policy('STARUSER', 'W_ACTIVITY_HISTORY_F',
'W_ACTIVITY_HISTORY_F_PC_POLICY');
    commit;
end;
/
```

- 2) Exchange the partition into a new table by running this command:

```
create table r1_p1 as select * from staruser.w_activity_history_f where
0=1;
alter table staruser.w_activity_history_f exchange subpartition r1_p1
with table r1_p1;
```

- 3) Export the table or move it to another database. See http://docs.oracle.com/cd/B19306_01/server.102/b14215/dp_export.htm for more information.

- 4) Remove the table from the database:

```
drop table r1_p1;
```

- 5) Restore security by running the staretl process for your operating system.

- ▶ In Windows, run:


```
staretl.bat "-from 31 -to 31"
```
- ▶ In Linux, run:


```
./staretl.sh -from 31 -to 31
```

Restoring Archived History

To restore archived history to the database:

- 1) Drop security on the W_ACTIVITY_HISTORY_F by running this command:

```
begin
    dbms_ols.drop_policy('STARUSER', 'W_ACTIVITY_HISTORY_F',
'W_ACTIVITY_HISTORY_F_P_POLICY');
    dbms_ols.drop_policy('STARUSER', 'W_ACTIVITY_HISTORY_F',
'W_ACTIVITY_HISTORY_F_PC_POLICY');
    commit;
end;
/
```

- 2) Import the table into the STARUSER database. See http://docs.oracle.com/cd/E11882_01/server.112/e10819/expimp.htm for more information.

- 3) Restore the archived table to the STARUSER database.

- 4) Restore the partition by exchanging the archive table into the correct partition. To restore the partition run this command:

```
alter table w_activity_history_f exchange subpartition r1_p1 with table
r1_p1;
```

- 5) Remove the table by running this command:

```
drop table r1_p1;
```

- 6) Restore security by running the staretl process for your operating system.

- ▶ In Windows, run:


```
staretl.bat "-from 31 -to 31"
```
- ▶ In Linux, run:


```
./staretl.sh -from 31 -to 31
```

Physical Components

You should consider three physical components for sizing the data warehouse:

- ▶ ETL Process Server
- ▶ Oracle ODS Database
- ▶ Oracle Star Database

These physical components will be treated as separate instances of the Oracle database, or a physical server, although this is not required. There is no direct size impact on the P6 EPPM Database beyond normal space usage, and it is not considered as one of the physical components to be sized.

ETL Process Server

While this server is the central controller of the ETL process, it represents only a small portion of physical space used. The only files, other than the ETL process files, are log files from each run and the Temporary Flat files for xlat (language translations), which are minimal.

Calculating Spread Sizes

The actual number of spread sizes depends on factors such as calendar work days, large differences in dates relative to the data date, ETL time duration, and the exclusion of zero value rows. A quick calculation uses a "best guess" on the average number of days for activities and resource assignments.

Total Activities: 1,000,000 X Average Activity Length: 5 = Total Spreads:
5,000,000

Queries for Spread Sizing

Activity Spread Estimate Based on Median Activity Length

```
select
median(
greatest(nvl(target_end_date,to_date('12122000','mmddyyyy'))
,nvl(act_end_date,to_date('12122000','mmddyyyy'))
,nvl(reend_date,to_date('12122000','mmddyyyy'))
,nvl(rem_late_end_date,to_date('12122000','mmddyyyy')) )
-
least(nvl(target_start_date,to_date('12122199','mmddyyyy'))
,nvl(act_start_date,to_date('12122199','mmddyyyy'))
,nvl(restart_date,to_date('12122199','mmddyyyy'))
,nvl(rem_late_start_date,to_date('12122199','mmddyyyy')) )
) * count(*) Spread_Rows
```

```
from task t inner join project p on p.proj_id = t.proj_id and orig_proj_id
is null
```

```
where task_type in ('TT_Task','TT_Rsrc')
```

Resource Assignment Spread Estimate Based on Median Activity Length

```
select
```

```
median(
```

```
greatest(nvl(tr.target_end_date,to_date('12122000','mmddyyyy'))
```

```
,nvl(tr.act_end_date,to_date('12122000','mmddyyyy'))
```

```
,nvl(tr.reend_date,to_date('12122000','mmddyyyy'))
```

```
,nvl(tr.rem_late_end_date,to_date('12122000','mmddyyyy')) )
```

```
-
```

```
least(nvl(tr.target_start_date,to_date('12122199','mmddyyyy'))
```

```
,nvl(tr.act_start_date,to_date('12122199','mmddyyyy'))
```

```
,nvl(tr.restart_date,to_date('12122199','mmddyyyy'))
```

```
,nvl(tr.rem_late_start_date,to_date('12122199','mmddyyyy')) )
```

```
) * count(*) Spread_Rows
```

```
from tasksrc tr inner join project p on p.proj_id = tr.proj_id and
orig_proj_id is null
```

```
inner join task t on t.task_id = tr.task_id
```

```
where task_type in ('TT_Task','TT_Rsrc')
```

Oracle ODS Database

The Oracle ODS database is an optional target database for operational-level reporting. It creates the views of the P6 EPPM base tables and their corresponding spread data tables. The space usage for the Oracle ODS Database can be derived from a combination of the size of the PMDB database and the size of the spread data. The Oracle ODS Database has the following types of table data:

- ▶ **Spread Tables** This is a combination of the detailed, daily spread data and aggregate tables.
- ▶ **Hierarchy Tables** These tables map the underlying hierarchical relationships (such as EPS, WBS, etc.).

Indexing in the Oracle ODS Database defaults to the same indexing as the P6 EPPM Database. This should be augmented and adjusted based on site-specific reporting needs.

Estimating the Size of the Star Database

Only the fact tables will be considered for Star Database sizing purposes because they are responsible for most of the data. Of the four fact tables in Star, two of the tables (W_ACTIVITY_SPREAD_F and W_RESOURCE_ASSIGNMENT_SPREAD_F) are identical to the equivalent spread tables in ODS (see **Estimating the Size of the ODS Database** (on page 22) for details). The rows for the remaining fact tables were calculated in the previous sections. Spread and resource limit data are initially loaded into holding tables (_FS suffix), so sizes are doubled for these tables.

Data Component	Calculation	Rows	Size Example
W_ACTIVITYSPREAD_F	300 bytes/row	5,000,000 x 2	3.0 GB
W_RESOURCE_ASSIGNMENT_SPREAD_F	175 bytes/row	5,000,000 x 2	1.6 GB
W_RESOURCE_LIMIT_F	70 bytes/row	1,825,000 x 2	0.125 GB
Dimensional and Temporary	20% of Spread	n/a	0.9 GB
W_Project_History_F	265 bytes/row	6000	0.04 GB
Total			5.6 GB

Note: The above sizing does not include the W_Activity_History_F and W_WBS_History_F tables, which will be the largest if the History Interval is set at Weekly and the History Level is set at Activity.

Estimating the Size of the ODS Database

The estimated size of the ODS Database is derived from two primary components: the estimated size of PMDB project data and the estimated size of the spread data. The majority of Oracle ODS Database data is simply a copy of the P6 EPPM Database tables. The remaining space usage comes mostly from the ActivitySpread and ResourceAssignmentSpread tables. Any remaining data will be estimated as a percentage of the spread data (including aggregate spread tables and hierarchies).

Data Component	Calculation	Rows	Size Example
ActivitySpread	300 bytes/row	5,000,000 x 2*	3.0 GB

Data Component	Calculation	Rows	Size Example
ResourceSpread	175 bytes/row	5,000,000 x 2*	1.6 GB
Other	30% of Spreads	n/a	0.4 GB
Total			45 GB

* Temporary data stored during loading process results in copies of spread data during initial ETL.

Oracle Star Database

The Oracle Star Database contains a dimensional data model along with several fact tables (including UDF fact tables and history fact tables (if configured)) and the supporting dimensions. Excluding the configurable history data tables, the Star Database will contain fewer rows because baseline projects are not directly accessible. The Oracle Star database is still much smaller than Oracle ODS Database.

The fact data represents the largest portion of data in the Oracle Star Database. As with any Star schema, this data is the most granular data built to support roll-up queries.

Depending on the History Interval and History Level settings, the Activity, WBS history tables and Slowly Changing Dimension (_HD) tables are likely to be the largest.

The primary two fact tables contain activity and resource assignment spread data respectively. The size of these tables will be the same as the corresponding Oracle ODS Database tables (ActivitySpread and ResourceAssignmentSpread).

The next largest fact table contains resource utilization data (W_RESOURCE_LIMIT_F). This differs from other fact tables in that the data size is not a function of the number or size of projects. Instead, it is a function of the number of resources in the database, and the size of the data warehouse reporting window. There is a daily value for everyday of the reporting period, and for each resource. For example, if the reporting window spans five years (1,825 days), and there are 1,000 resources in the database, the total recorded in the fact table will be 1,825,000.

The Project History fact table is the smallest and it has only project-level data. The difference is that this table is a trending table with snapshots of the data over time. The number of snapshots depends on the interval chosen during installation (weekly, monthly, financial period). The granularity of this fact table goes to only the project-level; it contains no spread information. Calculate the number of rows using the total non-baseline projects times the number of snapshots. This will grow over time, so the yearly total for a 10,000 project database with weekly snapshots will be 520,000 rows.

Physical Hardware

When evaluating the physical hardware requirements, you need to consider:

- ▶ The performance of the ETL process
- ▶ The performance and concurrency of the online reporting solution

While the ETL process is fixed regarding concurrency, the reporting needs will vary greatly. The demands on P6 Reporting Database and P6 Analytics may change from day-to-day. For performance sizing of OBI, please refer to the technical documents for the specific component (BI Publisher, BI Server, and Answers/Dashboards). This document will focus on the performance of the ETL process and queries generated against the warehouse databases (ODS and Star).

Monitoring Dimension and History Level Impacts on the Database

P6 Reporting Database includes slowly changing dimensions and daily level history for activity and resource assignment. If you choose to use slowly changing dimensions and these levels of history, they could cause the Star database to grow quickly. For example, if you update 10 activities in a project and run the ETL process, 10 rows will be added in the activity slowly changing dimension table and in the assignment history and activity history fact tables.

You should plan the database and server configuration to ensure you can monitor database growth and allow for expansion based on the number of changes. If growth is exceeding what you expected in your environment, you can disable slowly changing dimensions if they are not required. See the installation and configuration guides for more information.

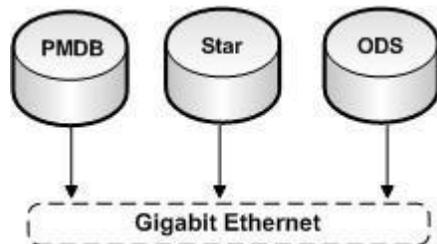
Consideration for ETL Performance

The ETL process for P6 Reporting Database was designed for multi-core processor systems. Instead of using a serialized process, Java is used to create a multi-threaded process to run concurrent threads of SQL*Plus, SQL*Loader. This can result in multiple run threads on the various servers. This also means that the process can be adversely affected by having to compete with other applications sharing the same resources. Therefore, an ideal configuration would have dedicated cores available for each component in the process and a dedicated database instance with sufficient SGA allocation and high through put IO system.

This is an ideal configuration that is meant to minimize contention. By dedicating resources to each of the physical components, concurrent performance will be maximized during peak usage. Different steps in the ETL process put a variety of load on each component. In this release, there is no concurrent processing occurring simultaneously on both ODS and Star servers. Therefore, from the ETL perspective, ODS and Star could share the same physical hardware. You should consider concurrent reporting usage when determining the correct CPU requirements for ODS and Star.

Network

While there is an advantage to separating components, doing so assumes that the network connections between servers have high bandwidth and low latency. These servers should be on the same local area network with gigabit connections. Any increase in latency will have an effect on the ETL performance.



The Star schema supports multiple P6 EPPM data sources for a single data warehouse instance. With the internationalization of the corporate operations, it's likely that the data sources might be distributed across geographic regions and connected via corporate WAN instead of the high-speed, low-latency LAN. Data extraction is a high volume operation, and the bandwidth and network latency impact performance. Oracle has a 10-10 rule for the corporate WAN connections:

- ▶ Minimum network bandwidth: 10 mbps
- ▶ Maximum network latency: 10 ms

For a successful deployment, you should have network parameters that are better than these minimum requirements.

Memory

The large number of parallel processes running on large sets of data will significantly impact the demands on memory. You should run the components of the data warehouse system on 64-bit operating systems to allow for large memory allocations. Constraining memory quickly reduces performance.

The database servers need block buffer and individual process memory. These servers should always be setup using Dedicated Server (not Shared Server). For an Oracle 11g database, Oracle recommends minimum MEMORY_TARGET of 2 GB (for 10g, set SGA_TARGET to the same minimum value). Otherwise, let the database server manage its own memory.

The Java process on the ETL Process Server is running multiple threads in the same process. Only run with a 64-bit version of the JRE to allow for larger memory allocation. The maximum memory allocation for the Java process is configurable during setup (Max Heap Size). The default is 1 GB. This may be inadequate for many datasets and may cause failures in the ETLCalc process. Start with a minimum of 4 GB of memory for the Java process.

Extract and Load

Customized ETL processes are implemented for the data extraction and loading. These processes are completely PL/SQL based and use only resources on the database server. Multiple database threads can run PL/SQL anytime. These bulk PL/SQL inserts, updates, and deletes make the process resource intensive. Oracle recommends 10 or more cores on the database server, so each thread has an available CPU to run. The performance of each core will determine the speed (rows/second) of the extract and load processes.

Summary of Physical Hardware Sizing

Use the guidelines in this section when planning for the physical hardware for the P6 Reporting Database.

Size of P6 EPPM Database

The size of the database will affect performance and the performance of the ETL process, since all records must be processed. The size of the database can also have some impact on the ETL process, since more project data will likely translate into more usage and more records in the database.

Amount of Change

You should monitor the performance of the Px services. While the database size plays a role, the volume of changes made will affect performance most. Even a small user community can generate a large amount of changes with functions like copy/paste, delete project, and create baseline, as well as simple changes to global resources like calendars. You should monitor the amount of changes made prior to installing the products to better plan the implementation.

In this version of the P6 Reporting Database, ODS and Star are always full ETL processes. Fresh data is pulled from the P6 EPPM schema and existing history table data are preserved during the ETL process. You must decide whether to implement a history data purge process.

Complexity of Project Data

Consider the data in the P6 EPPM Database. A very complex EPS/WBS structure will increase the processing time because these hierarchical elements require global processing even when the changes are small. The average size of activities can increase the memory requirements of the calculation process. Larger projects will cause more volume of change because the entire project is recalculated based on any changes within the project. This does not include changes to non-scheduling data, such as activity codes and UDFs.

Expectations and Operations

You may expect the ETL to complete in a smaller time frame for operational purposes. This expectation could increase the hardware requirements. The following sections detail requirements for each database size.

Small Database Deployment

ETL Server, Oracle ODS Database, and Oracle Star Database on same box:

CPU	Quad Core 3.46 GHz or equivalent
RAM	4 GB without Star; 8 GB with Star

Notes:

- The above requirements assume light Business Analytics (Oracle Star Database) usage. Otherwise, you should allocate a separate server for the Oracle Star Database.
- You should add 4 GB RAM if you installed OBIEE on the same box.

Medium Database Deployment

Oracle ODS Database:

CPU	Quad Core 3.46 GHz or equivalent
RAM	4 GB

ETL Server, Oracle STAR Database:

CPU	Quad Core 3.46 GHz or equivalent
RAM	8 GB

Large Database Deployment

ETL Server:

CPU	Quad Core 3.46 GHz or equivalent
RAM	8 GB

Oracle ODS Database:

CPU	8 Core 3.46 GHz or equivalent
RAM	16 GB

Oracle Star Database:

CPU	8 Core 3.46 GHz or equivalent
RAM	16 GB

Database Servers

	Oracle ODS Database	Oracle Star Database
TableSpace Autoextent	Yes	Yes
Recommended Temp Tablespace	Minimum 2 files with file size set at OS max	Minimum 2 files with file size set at OS max
Recommended UNDO Tablespace	Minimum 2 files with file size set at OS max	Minimum 2 files with file size set at OS max
Minimum Space	2.5 X P6 EPPM Database	Same as P6 EPPM Database
SGA	8GB	8GB

Planning Revisited

Timing Estimation

The following timing numbers are only an approximation. The actual numbers for your deployment depend on the data structure and the hardware used.

Initial ETL

	Number of Objects		
	Small	Medium	Large
Projects	200	1,000	15,000
Activities	100,000	500,000	4,000,000
Resources	5,000	5,000	5,000
Resource Assignments	30,000	100,000	2,000,000
Codes Mapped	1	1	200
UDFs Mapped	0	0	200
Project with History (percentage of total)*	5.00%	5.00%	50.00%

	Number of Objects		
	Small	Medium	Large
Backlog of ETL runs	--	--	--
STAR_DATx cumulative size**	3.5 GB	8 GB	90 GB
STAR_HSTx cumulative size**	< 1 GB	< 1 GB	< 1 GB
Full ETL Approximate completion time	< 1 hr	< 1 hr	~ 7 hours

Incremental ETL - 2 months of history and Slowly Changing Dimensions

	Number of Objects		
	Small	Medium	Large
Projects	200	1,000	15,000
Activities	100,000	500,000	4,000,000
Resources	5,000	5,000	5,000
Resource Assignments	30,000	100,000	2,000,000
Codes Mapped	1	1	200
UDFs Mapped	0	0	200
Project with History (percentage of total)*	5.00%	5.00%	50.00%
Backlog of ETL runs	60 days	60 days	60 days
STAR_DATx cumulative size	8 GB	17 GB	255 GB

	Number of Objects		
	Small	Medium	Large
STAR_HSTx cumulative size	< 1 GB	< 1 GB	60 GB
Full ETL Approximate completion time	~ 2 hrs	~ 3 hrs	~ 11 hours

*In performance runs, history level is all Activity Levels. This is the worst case scenario.

**Oracle data files: depending on how you manage your tablespace, you can have a different number of data files.

Factors Affecting ETL Runtime

All the following can affect ETL runtime.

Projects Published The number of projects and associated project data will have a direct correlation to the duration of the ETL run time. Carefully consider which projects require loading to the STAR database by filtering for these projects. Some considerations may be active projects or projects that meet some specific criteria by using a project filter.

Spread Interval In P6 EPPM, when configuring the service settings, the administrator must define the spread interval for each project. It is defined as a starting date to the current date plus a rolling interval. The definition of this interval has a direct correlation to the amount of spread data in Star: the larger the interval, the greater the number of spread records for each project. Oracle recommends keeping the rolling interval to a couple years from the current data.

Traditional History This feature keeps history at a level defined within a project at a specific interval. The number of projects that meet the designated history level and the intervals that have been crossed throughout the life cycle of the Star database will determine the size of these history tables. You should consider how features can optimize performance, such as:

- ▶ Selectively determining the number of projects that require history.
- ▶ Determining the level of storage granularity for each project with history. The higher the granularity, the fewer rows there will be over time.
- ▶ Turning off project history over time as projects end or close to minimize the cost of writing records for projects that aren't active anymore.
- ▶ Noting requirements during the initial configuration to correctly size the number of partitions for the Star database and the months to include in each partition. An even distribution of data per partition is optimal to increase performance.

Activity Level History & Slowly Changing Data In P6 Reporting Database, Activity Level History has been coupled with features that can impact performance of the ETL process. The most import of these features are Slowly Changing Data. When a project has Activity Level History enabled, changes are tracked at the field level and new dimension and fact records are created with each change. You should limit the projects with this history setting to ones that need this level of granularity and require tracking of changed data over time. The more projects with this setting, the more records will be stored in the activity history fact table and the slowly changing dimensions and fact tables.

Burn Down and Work Planning These two features are driven by the project having Activity Level History and a set of predefined User Defined Fields with appropriate values. The calculation of these projects with this feature enabled can be costly, so you should minimize the number of projects with these features.

The Number of User Defined Fields and Codes mapped During the initial configuration, select User Defined Fields and Codes to include within the Star database. Note the requirements to determine which User Defined Fields and Codes are needed in Star. The more UDFs and codes that are mapped, the more they'll affect the ETL process.

Steps in ETL that have the most significant run time due to settings:

ETL Step	Definition	Affecting Factors	Considerations
ActivityResourceAssignmentSpreads	Load activity and resource assignment spreads (fact_load)	<ul style="list-style-type: none"> ▶ number of projects published ▶ spread interval defined ▶ number of projects with activity level history ▶ changes to project data with activity level history over time 	<ul style="list-style-type: none"> ▶ using project filter ▶ reducing spread interval ▶ evaluating projects that need activity level history to either turn off or move to different history setting
MergeProjectHistory	Merge Project History (hist_merge)	<ul style="list-style-type: none"> ▶ number of projects with history defined ▶ size of existing history ▶ history intervals 	<ul style="list-style-type: none"> ▶ limiting projects with history ▶ increasing interval of history
SourceExtract	Extract data from source database (source_extract)	<ul style="list-style-type: none"> ▶ amount of data published 	<ul style="list-style-type: none"> ▶ using project filter ▶ thread count increase

ETL Step	Definition	Affecting Factors	Considerations
DimensionLoad	Load Dimension Tables (dim_load)	<ul style="list-style-type: none"> ▶ amount of data published ▶ projects with activity level history 	<ul style="list-style-type: none"> ▶ using project filter ▶ re-evaluation project with activity level history

Planning for P6 EPPM's Impact on Star

You have many different configuration options you can deploy to accommodate your environment. The configuration that works best will depend on the data size, data requirements, environment, and servers available.

The data flow from P6 into the Star database can affect your data warehouse.

Choosing Between Standard Edition and Enterprise Edition

The following table shows features supported with either the Standard Edition of P6 Reporting Database or the Enterprise Edition. If you are using a medium or large database, Oracle recommends using the Enterprise Edition.

Feature	Standard Edition	Enterprise Edition	Enterprise Edition with Partitioning
Multiple Datasources			X
Activity Level History			X
WBS Level History	X	X	X
Assignment Level History			X
Project Level History	X	X	X
Burndown Subject Area			X
Work Planning Subject Area			X

Feature	Standard Edition	Enterprise Edition	Enterprise Edition with Partitioning
Small Database Size	X	X	X
Medium Database Size		X	X
Large Database Size		X	X
Row Level Security		X	X

P6 Extended Schema's Affect on Star

The P6 Extended Schema is the primary source of data for the Star database. You can schedule services to continuously update data to allow for near-real time reporting, which allows the database to cover more data. The amount of data populated in the Extended Schema will have a cascading affect on the size of the Star database and the duration of the ETL process.

Date Range

In P6, you can define the date range for the Extended Schema. The range begins with the implementation date and a determined date range (for example, 3 years). Extending the range too far into the future will create large spread ranges. If you use the default range of 3 years, P6 will create a rolling window. As the services run, the window will continue to move. For example, January 1, 2012 and January 1, 2015 will become January 2, 2012 and January 2, 2015. If you have projects that last for three years, the projects will naturally fall in line and give you the spread data for this time frame.

Spread data represents a row in the database for each activity that has work on that day and for every activity in every project that is published. For example, if you have a project that has data reaching out 10 years and has 100 activities, you will have a spread row for 100 activities x 365 days per year x 10 years. Using this formula, data can add up quickly. The Extended Schema handles this and the Star ETL quickly moves this data.

If you do not need these projects in Star for Analytical reporting, you can still use them in the Extended Schema for reporting. You can use a Star filter to decide which projects to include in the ETL process. By default, all projects with the **Enable Publication** setting in the Extended Schema will transfer to Star.

Enable Publication

If you want to enable publication for projects, you must select the **Enable Publication** setting in P6. When you enable publication the Project Service will retrieve and calculate this project. The project will then be part of the Extended Schema. The Project Service will monitor how many changes have been made on the project and how long since the project was published.

Each project service run will write the number of changes on that project to a table. If it crosses this threshold, the project will queue for processing. If this project has any changes, it will be processed. Star will use the Enable Publication setting to determine which projects are pulled into the Star database.

History Settings on Project Level

If using P6 EPPM 8.2 SP1 or later, the project history level settings are available in P6 and can be set project by project; history will populate up from activity to WBS to project. If you set history at the Activity level, P6 will also populate WBS and Project history. You can set history for month, financial period, quarter, or year. Selecting Activity-level history means you can capture daily activity and resource assignment data regardless of the periodic setting in P6.

Define history levels selectively. History will impact how much data and what level the history fact tables in Star will process. Depending on your database size, enabling Activity level history for all projects could cause a long ETL process time. Turn on history only for those projects that require history. When projects end, turn off history settings to prevent performance issues.

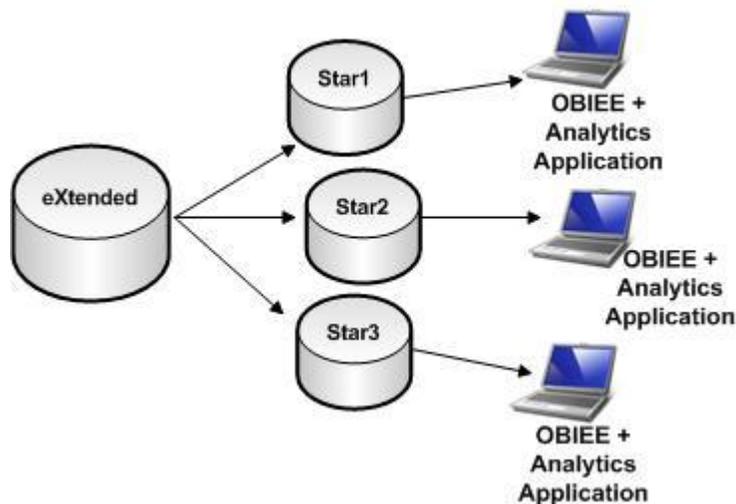
History is captured and rolled over in a merge to avoid overwriting existing history. This capture and merge can cause the ETL process to slow slightly after several historic period captures until crossing over a new partition. Partitioning is available to help performance and move older history into a new partition and help decrease ETL run time. Partition interval ranges are predefined in tables and index scripts in the Star installation\scripts directory. See the *P6 Analytics and Star Database Installation and Configuration Guide* for more information.

Slowly Changing Dimensions

Projects that have HIST_LEVEL HL_Task will be part of Slowly Changing Dimensions.

Multiple Star Environments

If you require low granularity (Activity, Daily) for a large number of projects, you may want to use multiple Stars. Multiple Stars will help decrease the run time and separate data into different ETL processes. To separate projects, you need to add a defining characteristic (portfolio, project code assignment, etc.) to the project in P6. The following diagram shows how the different Star environments will connect to their own OBI environment.



To set filters for the Star instance:

- 1) Follow the installation instructions in the *P6 Analytics and Star Database Installation and Configuration Guide*.
- 2) Define your filters as described in the *P6 Analytics and Star Database Installation and Configuration Guide*.

The filter will set a Select statement to retrieve the project ID's for only the projects you want in this Star. This can be any criteria discussed above such as portfolio or project code assignment. This filter mechanism overrides the default process of gathering all the project ID's for each project in the Extended Schema that has Enable Publication set.

Creating Multiple Data Sources for a Single Star Database

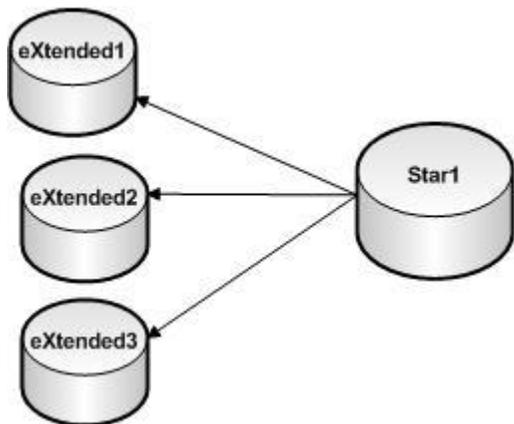
You can combine more than one Extended Schema into a single Star data warehouse, which allows you to have more than one P6 installation with an Extended Schema and be able to view all that data in one Star. For example, if you have a resource, John Smith, working in P6 database 1, you could see a combination of John Smith's assignments across database 1, 2, and 3. To work, this process requires that the names of resources, projects, or EPS's match. You can ensure this works by storing the data in the same dimensions and facts and inserting rows with unique row ID's for each data source. This gives you the ability to report on data across all data sources or report on data specific to just one data source.

Note: You cannot run the ETLs for the data sources at the same time. Each additional data source will require a unique installation with its own configuration information. Each installation will have its own staretl. The staretl processes cannot run at the same time because the unified Star will still use the same internal staging tables. If multiple ETL processes were running at the same time, they could overwrite each other in the staging tables and cause failures or data corruption.

Setting up Multiple Data Sources

You can create up to 99 data sources. You would create each data source individually and connect the specific Extended Schema where it should pull data. If you go above three data sources, you must adjust the scripts to accommodate the extra data sources.

You must determine how many data sources you require before you begin running ETL processes. You can also apply a filter to each data source to pull certain data from the extended schema. You must then schedule the ETL processes to run so they don't impact each other. The resulting environment could look similar to the diagram below:



For more details on how to set up multiple data sources, refer to the *P6 Analytics and Star Database Installation and Configuration Guide*.

Filtering by Data Source in Analytics

When creating analyses in OBI, a filter can be applied where:

data source id = (data source number)

In OBI, you can find a data source ID field under Project – General in most subject areas. You can apply the data source ID to any Analysis and set a value to limit the output to only that desired data source. You will see a similar data source ID in resource-related Subject areas under Resource – General.

Single Star, Single Data Source, Multiple ETLs

In this case, you have one primary data source pulling from the Extended Schema (Px). There is only one Star data warehouse, but you will have different ETL processes. You may need different ETL processes to update certain projects for a different company division at a different time. Because the ETL process is typically meant to run once a day due to limitations of time and staging tables, you could use filtering and multiple data sources inside of one Star instance.

Caution: ETL processes cannot run at the same time. You cannot overlap ETL processes because the staging tables inside the Star are shared and any overlap can cause data corruption.

Setting up Multiple ETLs

To setup the first ETL:

- 1) Install the Reporting Database per instructions in the *P6 Analytics and Star Database Installation and Configuration Guide*.
- 2) Define a filter against Px (see the *P6 Analytics and Star Database Installation and Configuration Guide*). Set the filter value in your properties file for filter 1. This will be data source 1.
- 3) Run the **staretl** process by running **staretl.bat** or **staretl.sh**.
- 4) Ensure Star data is using this filter.

To setup the second ETL:

- 1) Rename the Reporting Database folder that was installed.

Note: You can have folders with the same name or you can choose a new name for the folder when you install the second ETL for P6 Reporting Database. Choosing a new name will ensure the installer doesn't try to rename the folder after the install is done.

- 2) Install the Reporting Database as you usually would to a new location.
- 3) Define a new filter against Px (see the *P6 Analytics and Star Database Installation and Configuration Guide*). Set the filter value in your properties file for filter 2. This will be data source 2.
- 4) Create a second database link for Staruser.
- 5) Run the ETL process by executing **staretl.bat "-s2"** or **staretl.sh "-s2"**.
- 6) Ensure Star data with a data source ID of 2 is using this filter.

Run the ETLs

To setup the ETL:

- 1) Rename the folders for the first data source back to the original name.
- 2) Run the ETL from that folder by running **staretl.bat "-s1"** or the **staretl.sh "-s1"**.

- 3) Determine how long each ETL process takes. If the first ETL takes 4 hours and the second takes 6 hours, you need to set the jobs to run at separate times and allow extra time if one ETL run goes longer than usual. You should monitor this as ETL runtimes and Star history fact tables can grow as data grows in P6.
- 4) Run the ETL for the second data source as **staretl.bat “-s2”** or the **staretl.sh “-s2”**.

Tips

- ▶ When running the ETL, always use the parameters you have set for your data sources.
- ▶ In P6 Analytics this data will combine as if it were one Star until a filter is applied. Under **Project – General** and **Resource – General** there is a field for data source id that can be used to filter on either data source 1 or data source 2.
- ▶ If adding more than three data sources, you will need to manually alter the scripts (See the *P6 Analytics and Star Database Installation and Configuration Guide*). Post Analytics 3.1 release, unlimited data sources are supported.

Appendix: Checklists

Checklist for P6 Analytics and the Star Database

Review the details of this checklist before installing P6 Analytics. Use the checklist to ensure you successfully install P6 Analytics. The checklist contains details on environment settings you need to consider, such as memory and disk space.

Pre-installation for an Oracle Platform

Ensure you:

Check	Item
	Downloaded P6 Analytics and the Star database software for the correct operating system. The Star database is part of the Primavera P6 Reporting Database software. (See the Oracle Software Delivery Cloud.)
	Installed the correct JRE version (see the <i>Tested Configurations</i> document).
	Installed the correct JRE platform (x86 versus x64).
	Can connect to the P6 EPPM database from this machine.
	Installed P6 EPPM and verified which release you have. Note: See the <i>Tested Configurations</i> for supported versions.
	Have the same Oracle database version on all servers.

Check	Item
	<p>Enabled partitioning on the Star instance.</p> <p>Note: Oracle recommends you use partitioning if the P6 EPPM database size exceeds 'Small' as defined in the <i>P6 Analytics and P6 Reporting Database Planning and Sizing Guide</i> or if multiple P6 data sources are connected to the Star.</p>

Pre-installation for a Microsoft SQL Server Platform

Ensure you:

Check	Item
	Downloaded P6 Analytics and the Star database software for the correct operating system. The Star database is part of the Primavera P6 Reporting Database software. (See the Oracle Software Delivery Cloud.)
	Installed the correct JRE version (see the <i>Tested Configurations</i> document).
	Installed the correct JRE platform (x86 versus x64).
	Can connect to the P6 EPPM database from this machine.
	<p>Installed P6 EPPM and verified which release you have.</p> <p>Note: See the <i>Tested Configurations</i> for supported versions.</p>
	Installed the Oracle Gateway (10g or 11g).
	<p>Can connect the Oracle Gateway to the P6 EPPM database.</p> <p>Note: To determine whether you have a connection, execute a SELECT statement through Oracle Gateway.</p>

Pre-installation for Configuring the ETL Server

Ensure you:

Note: This section is based on large implementations. Refer to the ***Summary of Physical Hardware Sizing*** (on page 26) section for details on small and medium implementations.

Check	Item
	Reviewed the <i>P6 Analytics and P6 Reporting Database Planning and Sizing Guide</i> for sizing and spacing guidelines.
	<p>Can connect to the P6 EPPM database.</p> <p>Note: If you are using Microsoft SQL Server, ensure you can connect to the P6 EPPM database through the Oracle Gateway.</p>

Check	Item
	Can connect to the Star database instance.
	Configured TNSNAMES correctly with all instances.
	Defined the ORACLE_HOME path on the server running the ETL scripts so it can find SQLLDR.
	Can connect JDBC and OCI on the ETL server. Note: You can test the connection with a sample Java application.
	Verified the user running scripts on the server has Read, Write permissions to install the directory. Note: On UNIX systems, create and install the database as 'Oracle' user.
	Placed all servers in same data center with a Gigabit Ethernet connection between servers.
	Verified each ETL Process server has 8-12 core processors.
	Verified each database server has 8-10 core processors.
	Are running the ETL Server on 64-bit machine to allow for large memory allocations.
	Set the Java Max Heap setting to at least 4 GB of memory. Allow for more memory, if available, for the Java process.
	Set the Bulk Load files location from the ETL process to a location with a large amount of available disk space. Note: These files can become large depending on the size of the database. Available space should range from a few gigs to many tens of gigabytes.
	Are using a supported Operating System. See the <i>Tested Configurations</i> document for Supported Operating Systems.

Pre-installation for the Star Database

Ensure you:

Check	Item
	Reviewed the <i>P6 Analytics and P6 Reporting Database Planning and Sizing Guide</i> for sizing and spacing guidelines.
	Have a tablespace called STAR_DAT1.
	Have a tablespace called STAR_HST1.
	Set the tablespace with Auto Extend On.
	Set the tablespace extent management to Local.

Check	Item
	Set the character set to UTF8 or WE8MSWIN1252. P6 Analytics also supports the AL32UTF8 character set.
	Set the character set to the same character set used in the P6 EPPM database.
	Set the Star tablespace to at least the same size as the P6 EPPM database.
	Run the queries to validate database settings. See the database scripts listed in Appendix A2 at the end of this document.
	Set datafiles to autoextendable.
	Set maxbytes to the OS's maximum, which is typically 32 GB with minimum 2 files. Use 3 to 5 files for larger systems.
	Run the queries to validate the UNDO Tablespace settings. See the database scripts listed in Appendix A2 at the end of this document.
	Run the queries to validate the SGA settings on the database. See the database scripts listed in Appendix A2 at the end of this document. Note: Larger systems require 8 GB.
	Run the queries to validate the Archive Log Mode on the database. See the database scripts listed in Appendix A2 at the end of this document. Note: This should be in NON archive log mode.
	Execute the following SQL to verify the tablespace: <pre>select s.tablespace_name, f.file_name, f.maxbytes, f.autoextendable from dba_data_files f inner join dba_tablespaces s on s.tablespace_name = f.tablespace_name where s.tablespace_name like '%STAR%';</pre>
	Have both block buffer and individual process memory available.
	Setup the database instance using Dedicated Server (not Shared Server).
	Set the minimal value set for the MEMORY_TARGET variable: For an Oracle 11g database, the recommended minimum MEMORY_TARGET is 2 GB. For an Oracle 10g, set SGA_TARGET to the same minimum value. For all other versions, let the database server manage its own memory.
	Can connect JDBC and OCI on the Star server. Note: You can test the connection with a sample Java application.
	Decided whether there will be multiple data sources. The sizing requirements will apply for each individual data source included.
	Verified that you are using one of the following database versions. See the Tested Configurations document for more information on supported databases.

Installation: Installer And Configuration

Ensure you:

Check	Item
	Reviewed the <i>P6 Analytics and Star Database Installation and Configuration Guide</i> .
	Set the JAVA_HOME and PATH definition. Note: The config.cmd(.sh) needs JAVA_HOME set in order to launch. For example, in Windows, set JAVA_HOME=C:\Program Files\Java\jre1.7.0_40.
	Determine your date range for your data warehouse. Note: Determine your date range before running the installer. This date range will include spread data and the start date and rolling date range of your data warehouse. See the <i>Installation and Configuration Guide</i> for more information.
	Determine a reasonable date range setting. Note: If you set a large data range this will result in a very large amount of spread data. A large amount of spread data requires more database space, more hard drive space for temp files, and increased run times. Consider your date range carefully. Spreads outside of the date range will combine into one bucket at the beginning and end of the date range.
	Determine your Project Trend Interval. Note: Project Trend Interval will determine how to store your project history. See the <i>Installation and Configuration Guide</i> for more information.
	Determine which activity, project, and resource codes you will use before installation. Note: Dynamic codes are available for slicing data in OBI. Code hierarchies are also supported if you select them when you install this application. <ul style="list-style-type: none"> ▶ Matching Criteria/Regular expression: The name of the code as registered in Project Management. ▶ Name: The display name for the code name label in OBI. ▶ Description: The display name for the code value label in OBI.
	Determine which activity, project, and resource UDFS you will use before installation. <ul style="list-style-type: none"> ▶ Matching Criteria/Regular expression: The name of the UDF as registered in P6. ▶ Name: The display name for the UDF name label in OBI. ▶ Description: The display name for the UDF value label in OBI.
	Consider the available disk space when you set the Logging level. Note: If DEBUG is set, the logs can become large. Oracle recommends INFO for a logging level unless troubleshooting is an issue.

Check	Item
	<p>Set the Java Max Heap setting to at least 4 GB of memory. Allow for more memory, if available, for the Java process.</p> <p>Note: The default for the Maximum Java Heap Size (MB) parameter is 1,204MB. The minimum value is 512MB. The value assigned must be a multiple of 512MB. The <i>Planning and Sizing Guide</i> references 1 GB, but notes that 1 GB may be inadequate for many datasets and may cause failures in the ETLCalc process. Oracle recommends starting with a minimum of 4 GB of memory for the Java process.</p>
	<p>Have the P6 EPPM privuser username/password information available.</p>
	<p>Have the Star instance and System username/password information available.</p>
	<p>Determine which P6 EPPM users will be reporting users.</p> <p>Note: Reporting users will have a database user created for them. That means you must set the report_user_flag field to 'Y' on the USERS table in the P6 EPPM database for all those users who will become reporting users.</p>
	<p>Define Financial Periods before running the P6 Reporting Database processes.</p> <p>Note: If you will set the Project Trend intervals based on the P6 EPPM Financial Period, you must set the Financial Period before running the P6 Reporting Database processes.</p>
	<p>Set the Bulk Load files location for the ETL process to a location with a large amount of available disk space.</p> <p>Note: These files can become large depending on the size of the database. You need ranges from a few gigs to many tens of gigabytes of available space.</p>
	<p>Set Oracle database user requirements for users and passwords when creating the Oracle stageuser and staruser.</p>

Installation: Running The Processes

Ensure you:

Check	Item
	Grant EXECUTE to PUBLIC for the standard Oracle SYS.DBMS_RANDOM package to run the RunETL Process.
	Locate the installation home directory that contains the staretl.bat (or staretl.sh).
	View the log and html file for the ETL process in the installation home directory to view the status of the process and verify it completed successfully.

Post-installation: After The Processes Complete

Ensure you:

Check	Item
	Run the staretl.bat file.
	In the installation directory, review the etlprocess.log and etlprocess.html log to ensure processes completed successfully.
	Ensure the system created an staretl.bat and the appropriate log files.
	See the <i>P6 Analytics and Star Database Installation and Configuration Guide</i> for more information on logs.

Database Scripts For Validations

Use the following scripts to validate your database:

Script	Description
SGA (Memory) Validation	Requires 8 GB for larger systems. Script: SQL> SELECT * FROM v\$sgainfo; SQL> show sga
Check archive log mode	For ERDB, the database should be in non-archive mode. Script: SQL> SELECT log_mode FROM v\$database;

Script	Description
Processes Init Parameter	Requires a minimum of 300. Script: <pre>SQL> show parameter process</pre>
Temp Tablespace Validation	Requires a minimum of 2 files with file size set at the OS maximum, typically 32 GB. Larger systems should have 3 files for the temp tablespace. Script: <pre>SQL> select TABLESPACE_NAME, FILE_ID, BYTES_USED, BYTES_FREE from V\$TEMP_SPACE_HEADER;</pre>
UNDO Tablespace Validation	Set the datafile to be autoextensible. Set the maxbytes to the OS maximum, typically 32 GB. Requires a minimum of 2 files; 3 to 5 files for larger systems. Script: <pre>SQL> select s.tablespace_name, f.file_name, f.maxbytes, f.autoextensible from dba_data_files f inner join dba_tablespaces s on s.tablespace_name = f.tablespace_name where s.tablespace_name like '%UNDO%';</pre>
Tablespaces for Star	Set the datafile to be autoextensible. Set the maxbytes to the OS maximum, typically 32 GB. Requires a minimum of 2 files; 3 to 5 files for larger systems. Script: <pre>SQL> select s.tablespace_name, f.file_name, f.maxbytes, f.autoextensible from dba_data_files f inner join dba_tablespaces s on s.tablespace_name = f.tablespace_name where s.tablespace_name like '%STAR%';</pre>

Checklist for ODS Database

Review the details of this checklist before installing P6 Reporting Database for ODS. Use the checklist to ensure you successfully install P6 Reporting Database for ODS. The checklist contains details on environment settings you need to consider, such as memory and disk space.

Pre-installation for an Oracle Platform

Ensure you:

Check	Item
	Downloaded P6 Reporting Database software for the correct operating system. (See the Oracle Software Delivery Cloud.)
	Installed the correct JRE version (see the <i>Tested Configurations</i> document).
	Installed the correct JRE platform (x86 or x64) (see the <i>Tested Configurations</i> document).
	Can connect to the P6 EPPM database from this machine.
	Installed P6 EPPM and verified which release you have. Note: See the <i>Tested Configurations</i> for supported versions.
	Have the same Oracle database version on all servers.

Pre-installation for a Microsoft SQL Server Platform

Ensure you:

Check	Item
	Downloaded P6 Reporting Database for the correct operating system. (See the Oracle Software Delivery Cloud.)
	Installed the correct JRE version (see the <i>Tested Configurations</i> document).
	Installed the correct JRE platform (x86 versus x64).
	Can connect to the P6 EPPM database from this machine.
	Installed P6 EPPM and verified which release you have. Note: See the <i>Tested Configurations</i> for supported versions.
	Installed the Oracle Gateway (10g or 11g).
	Can connect the Oracle Gateway to the P6 EPPM database. Note: To determine whether you have a connection, execute a SELECT statement through Oracle Gateway.

Pre-installation for Configuring the ETL Server

Ensure you:

Check	Item
	Reviewed the <i>P6 Analytics and P6 Reporting Database Planning and Sizing Guide</i> for sizing and spacing guidelines.
	Can connect to the P6 EPPM database. Note: If you are using Microsoft SQL Server, ensure you can connect to the P6 EPPM database through the Oracle Gateway.
	Can connect to the ODS database instance.
	Configured TNSNAMES correctly with all instances.
	Defined the ORACLE_HOME path on the server running the ETL scripts so it can find SQLLDR.
	Can connect JDBC and OCI on the ETL server. Note: You can test the connection with a sample Java application.
	Verified the user running scripts on the server has Read, Write permissions to install the directory. Note: On UNIX systems, create and install the database as 'Oracle' user.
	Placed all servers in same data center with a Gigabit Ethernet connection between servers.
	Verified each ETL Process server has 8-12 core processors.
	Verified each database server has 8-10 core processors.
	Are running the ETL Server on 64-bit machine to allow for large memory allocations.
	Set the Java Max Heap setting to at least 4 GB of memory. Allow for more memory, if available, for the Java process.
	Set the Bulk Load files location from the ETL process to a location with a large amount of available disk space. Note: These files can become large depending on the size of the database. Available space should range from a few gigs to many tens of gigabytes.
	Are using a supported Operating System. See the <i>Tested Configurations</i> document for Supported Operating Systems.

Pre-installation for the ODS Database

Ensure you:

Check	Item
	Reviewed the <i>P6 Analytics and P6 Reporting Database Planning and Sizing Guide</i> for sizing and spacing guidelines.
	Have a tablespace called ODS_DAT1.
	Set the tablespace with Auto Extend On.
	Set the tablespace extent management to Local.
	Set the character set to UTF8 or WE8MSWIN1252. P6 Analytics also supports the AL32UTF8 character set.
	Set the character set to the same character set used in the P6 EPPM database.
	Set the ODS tablespace to at least 2.5 times larger than the size of the P6 EPPM database.
	Run the queries to validate database settings. See the database scripts listed in Appendix A2 at the end of this document.
	Set datafiles to autoextensible.
	Set maxbytes to the OS's maximum, which is typically 32 GB with minimum 2 files. Use 3 to 5 files for larger systems.
	Run the queries to validate the UNDO Tablespace settings. See the database scripts listed in Appendix A2 at the end of this document.
	Run the queries to validate the SGA settings on the database. See the database scripts listed in Appendix A2 at the end of this document. Note: Larger systems require 8 GB.
	Run the queries to validate the Archive Log Mode on the database. See the database scripts listed in Appendix A2 at the end of this document. Note: This should be in NON archive log mode.
	Execute the following SQL to verify the tablespace: <pre>select s.tablespace_name, f.file_name, f.maxbytes, f.autoextensible from dba_data_files f inner join dba_tablespaces s on s.tablespace_name = f.tablespace_name where s.tablespace_name like '%ODS%';</pre>
	Have both block buffer and individual process memory available.
	Setup the database instance using Dedicated Server (not Shared Server).

Check	Item
	Set the minimal value set for the MEMORY_TARGET variable: For an Oracle 11g database, the recommended minimum MEMORY_TARGET is 2 GB. For an Oracle 10g, set SGA_TARGET to the same minimum value. For all other versions, let the database server manage its own memory.
	Can connect JDBC and OCI on the ODS server. Note: You can test the connection with a sample Java application.
	Verified that you are using one of the following database versions. See the <i>Tested Configurations</i> document for more information on supported databases.

Installation: Installer And Configuration

Ensure you:

Check	Item
	Reviewed the <i>P6 Reporting Database for ODS Installation and Configuration Guide</i> .
	Set the JAVA_HOME and PATH definition. Note: The config.cmd(.sh) needs JAVA_HOME set in order to launch. For example, in Windows, set JAVA_HOME=C:\Program Files\Java\jre1.6.0_37
	Determine your date range for your data warehouse. Note: Determine your date range before running the installer. This date range will include spread data and the start date and rolling date range of your data warehouse. See the <i>P6 Reporting Database for ODS Installation and Configuration Guide</i> for more information.
	Determine a reasonable date range setting. Note: If you set a large data range this will result in a very large amount of spread data. A large amount of spread data requires more database space, more hard drive space for temp files, and increased run times. Consider your date range carefully. Spreads outside of the date range will combine into one bucket at the beginning and end of the date range.
	Consider the available disk space when you set the Logging level. Note: If DEBUG is set, the logs can become large. Oracle recommends INFO for a logging level unless troubleshooting is an issue.

Check	Item
	<p>Set the Java Max Heap setting to at least 4 GB of memory. Allow for more memory, if available, for the Java process.</p> <p>Note: The default for the Maximum Java Heap Size (MB) parameter is 1,204MB. The minimum value is 512MB. The value assigned must be a multiple of 512MB. The <i>Planning and Sizing Guide</i> references 1 GB, but notes that 1 GB may be inadequate for many datasets and may cause failures in the ETLCalc process. Oracle recommends starting with a minimum of 4 GB of memory for the Java process.</p>
	Have the P6 EPPM privuser username/password information available.
	Have the ODS instance and System username/password information available.
	<p>Determine which P6 EPPM users will be reporting users.</p> <p>Note: Reporting users will have a database user created for them. That means you must set the report_user_flag field to 'Y' on the USERS table in the P6 EPPM database for all those users who will become reporting users.</p>
	<p>Set the Bulk Load files location for the ETL process to a location with a large amount of available disk space.</p> <p>Note: These files can become large depending on the size of the database. You need ranges from a few gigs to many tens of gigabytes of available space.</p>
	Set Oracle database user requirements for users and passwords when creating the Oracle odsuser.

Installation: Running The Processes

Ensure you:

Check	Item
	Grant EXECUTE to PUBLIC for the standard Oracle SYS.DBMS_RANDOM package to run the RunETL Process.
	Locate the installation home directory that contains the odsetl.bat (or odsetl.sh).
	View the log and html file for the ETL process in the installation home directory to view the status of the process and verify it completed successfully.

Post-installation: After The Processes Complete

Ensure you:

Check	Item
	Run the odsetl.bat file.
	In the installation directory, review the etlprocess.log and etlprocess.html log to ensure processes completed successfully.
	Ensure the system created an odsetl.bat and the appropriate log files.
	See the <i>P6 Reporting Database for ODS Installation and Configuration Guide</i> for more information on logs.

Database Scripts For Validations

Use the following scripts to validate your database:

Script	Description
SGA (Memory) Validation	Requires 8 GB for larger systems. Script: SQL> SELECT * FROM v\$sgainfo; SQL> show sga
Check archive log mode	For ERDB, the database should be in non-archive mode. Script: SQL> SELECT log_mode FROM v\$database;
Processes Init Parameter	Requires a minimum of 300. Script: SQL> show parameter process
Temp Tablespace Validation	Requires a minimum of 2 files with file size set at the OS maximum, typically 32 GB. Larger systems should have 3 files for the temp tablespace. Script: SQL> select TABLESPACE_NAME, FILE_ID, BYTES_USED, BYTES_FREE from V\$TEMP_SPACE_HEADER;

Script	Description
UNDO Tablespace Validation	<p>Set the datafile to be autoextensible. Set the maxbytes to the OS maximum, typically 32 GB. Requires a minimum of 2 files; 3 to 5 files for larger systems.</p> <p>Script:</p> <pre>SQL> select s.tablespace_name, f.file_name, f.maxbytes, f.autoextensible from dba_data_files f inner join dba_tablespaces s on s.tablespace_name = f.tablespace_name where s.tablespace_name like '%UNDO%';</pre>
Tablespaces for ODS	<p>Set the datafile to be autoextensible. Set the maxbytes to the OS maximum, typically 32 GB. Requires a minimum of 2 files; 3 to 5 files for larger systems.</p> <p>Script:</p> <pre>SQL> select s.tablespace_name, f.file_name, f.maxbytes, f.autoextensible from dba_data_files f inner join dba_tablespaces s on s.tablespace_name = f.tablespace_name where s.tablespace_name like '%ODS%'</pre>

For More Information

In This Section

Where to Get Documentation	53
Where to Get Training	55
Where to Get Support	55

Where to Get Documentation

For the most up-to-date versions of all manuals and technical documents related to installing, administering, and using P6 Analytics, go to:

http://download.oracle.com/docs/cd/E49048_01/index.htm

Most documentation assumes a standard setup of the product, with full access rights to all features and functions.

You can also access the versions of the product manuals and technical documents that were available at the time of the release from the P6 Analytics Documentation Center, located in the \Documentation\Documentation_library\language folder of the P6 Analytics physical media or download.

The following table describes the core documents available for P6 Analytics and lists the recommended readers by role.

Title	Description
<i>What's New in P6 Analytics</i>	This guide highlights the new and enhanced features included in this release. You can also use the <i>Cumulative Feature Overview Tool</i> to identify the features that have been added since a specific release level. All users should read this guide.
<i>P6 Analytics and P6 Reporting Database Planning and Sizing Guide</i>	This guide details how to plan your installation and ensures you have the necessary technical specifications to successfully install P6 Analytics and P6 Reporting Database. It also includes checklists for P6 Analytics and P6 Reporting Database to help guide you through the installation. All administrators should read this guide.

Title	Description
<i>P6 Analytics and Star Database Installation and Configuration Guide</i>	<p>This guide gives step-by-step instructions for installing and configuring P6 Analytics and the Star database portion of P6 Reporting Database.</p> <p>All administrators should read this guide.</p>
<i>P6 Reporting Database for ODS Installation and Configuration Guide</i>	<p>This guide explains how to install and configure the ODS portion of P6 Reporting Database. It describes how to install and configure the Oracle Gateway if the P6 Reporting Database is installed on a Microsoft SQL Server. It also provides information about how to run the Configuration Utility and configure P6 Reporting Database with BI Publisher.</p> <p>All administrators should read this guide.</p>
<i>P6 Analytics Post Installation Administrator's Guide</i>	<p>This guide provides information about P6 Analytics administrative tasks. It also includes information for Star security configuration, OBI installation and configuration, Financial Periods installation and configuration, and for configuring the Secure Sockets layer.</p> <p>All administrators should read this guide.</p>
<i>P6 Analytics Reference Manual</i>	<p>This manual has examples of sample dashboards and Burn Down activity use cases. It also tells users how to get started with P6 Analytics.</p> <p>All non-administrator users should read this guide.</p>
<i>P6 EPPM and P6 Analytics 3.2 System Architecture Data Sheet</i>	<p>The data sheet provides information on P6 EPPM, P6 Analytics, and P6 Reporting Database. It also provides a diagram to show how all products work together.</p> <p>All administrators should read this guide.</p>
<i>Security Guidance for P6 Analytics and P6 Reporting Database</i>	<p>This guide enables you to plan your security strategy for P6 Analytics and P6 Reporting Database. It includes information on safe deployments, authentication options, and specific security settings for the Star and ODS database.</p> <p>All administrators should read this guide.</p>
<i>Tested Configurations</i>	<p>Lists the configurations that have been tested and verified to work with P6 Analytics.</p> <p>The network administrator/database administrator and P6 Analytics administrator should read this document.</p>

Distributing Information to the Team

You can copy the online documentation to a network drive for access by project participants. Each team member can then view or print those portions that specifically relate to his or her role in the organization.

Throughout this documentation, the Security Guidance icon  helps you to quickly identify security-related content to consider during the installation and configuration process.

Where to Get Training

To access comprehensive training for all Primavera products, go to:

<http://education.oracle.com>

Oracle Learning Library

The Oracle Learning Library (OLL) provides online learning content covering Primavera products. Content includes videos, tutorials, articles, demos, step-by-step instructions to accomplish specific tasks, and self-paced interactive learning modules.

To access the learning library's Primavera content, go to:

<http://www.oracle.com/goto/oll>

Where to Get Support

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/us/support/contact-068555.html> or visit <http://www.oracle.com/us/corporate/accessibility/support/index.html> if you are hearing impaired.

Using Primavera's Support Resource Centers

Primavera's Support Resource Center provides links to important support and product information. Primavera's Product Information Centers (PICs) organize documents found on My Oracle Support (MOS), providing quick access to product and version specific information such as important knowledge documents, Release Value Propositions, and Oracle University training. PICs also offer documentation on Lifetime Management, from planning to installs, upgrades, and maintenance.

Visit <https://support.oracle.com/epmos/faces/DocumentDisplay?id=1486951.1> to access links to all of the current PICs.

PICs also provide access to:

- ▶ **Communities** which are moderated by Oracle providing a place for collaboration among industry peers to share best practices.
- ▶ **News** from our development and strategy groups.
- ▶ **Education** via a list of available Primavera product trainings through Oracle University. The Oracle Advisor Webcast program brings interactive expertise straight to the desktop using Oracle Web Conferencing technology. This capability brings you and Oracle experts together to access information about support services, products, technologies, best practices, and more.

Legal Notices

Oracle Primavera P6 Analytics and P6 Reporting Database Planning and Sizing Guide

Copyright © 1999, 2013, Oracle and/or its affiliates. All rights reserved.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate failsafe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

This software or hardware and documentation may provide access to or information on content, products and services from third-parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.