

Essbase® Analytic Services Release 7.1

This document describes the new features in Essbase Analytic Services (Release 7.0 and Release 7.1). For information on the new features in Administration Services, see the *Essbase Administration Services New Features* booklet.

This document contains the following information:

- [“New Features in Release 7.1” on page 3](#)
- [“New Features in Release 7.0” on page 9](#)

Contents

New Features in Release 7.1	3
Aggregate Storage Databases	3
Hyperion Remote Authentication Module	4
Unicode	4
Triggers on Data Changes	5
Reduced Security-File Fragmentation	6
MaxL	6
API	7
Platform Support Changes	8
Documentation	8
New Features in Release 7.0	9
MetaData Security	9
User Authentication Through the Hyperion Security Platform	10
MaxL	10
Report Writer	15
Performance	16
API	17
Calculation	18
Dimension Build	19
Metadata Sampling	20
Unicode Support	21
Data Mining Support	23
Monitoring Data Changes: Triggers	25
Administration	25
Delayed Free Space Recovery	27
Server Threads Specified by Application	28
Attribute Support for Hybrid Analysis Members	29
Platform Support Changes	29
Documentation	30
Analytic Services Documentation Information Map	30
Where to Get More Information	32

New Features in Release 7.1

The new features in Release 7.1 are described in the following topics:

- [“Aggregate Storage Databases” on page 3](#)
- [“Hyperion Remote Authentication Module” on page 4](#)
- [“Unicode” on page 4](#)
- [“Triggers on Data Changes” on page 5](#)
- [“Reduced Security-File Fragmentation” on page 6](#)
- [“MaxL” on page 6](#)
- [“API” on page 7](#)
- [“Platform Support Changes” on page 8](#)
- [“Documentation” on page 8](#)

Aggregate Storage Databases

Support for aggregate storage databases (provided by a new aggregate storage kernel) enables dramatic scalability improvements in both database aggregation time and dimensional scalability. The aggregate storage kernel does not replace the current block storage kernel. Aggregate storage databases typically address read-only, “rack and stack” applications that have large dimensionality, such as the following applications:

- Customer analysis—data is analyzed from any dimension and there are potentially millions of customers
- Procurement analysis—many products are tracked across many vendors
- Logistics analysis—near real-time updates of product shipments

Aggregate storage databases with many dimensions calculate (aggregate) in minutes.

For additional information about aggregate storage databases, see the *Essbase Analytic Services Database Administrator’s Guide*.

Outline Paging for Aggregate Storage Databases

Aggregate storage databases may have outlines with millions of members. Support for outline paging significantly reduces memory usage, enhancing the end-user experience. Depending on the operating system and available memory, you can create outlines with up to 20 million members.

The following new settings and commands control the memory available and provide information on the memory usage:

- OPGCACHESIZE
- GETOPGCACHESTATISTICS
- GETOPGDIMSTATISTICS

For additional information about this feature, see the *Essbase Analytic Services Database Administrator's Guide*. For specific details about the settings and commands, see the *Technical Reference*.

Hyperion Remote Authentication Module

The Hyperion Remote Authentication Module is an optional component of Hyperion security platform. Hyperion security platform is a framework that provides Hyperion applications with external authentication and single sign-on capabilities. For more information on Hyperion Security Platform, see [“User Authentication Through the Hyperion Security Platform” on page 10](#).

The Hyperion Remote Authentication Module is applicable to the following circumstances:

- Hyperion products are deployed in a UNIX environment with authentication occurring against a Windows domain.
- Hyperion products are deployed in a Windows environment, and users must log in to more than one Windows domain without changing the trust relationships.

You can obtain the Hyperion Remote Authentication Module from the Hyperion Download Center or from the Hyperion Hub CD.

Unicode

Unicode support now extends to the following features:

- **Triggers.** Triggers can be set for databases in Unicode-mode applications. For Unicode-mode databases, the spool file and the `.trg` file, which contains the trigger specifications, are encoded in UTF-8. E-mail notifications for triggers are in UTF-8 and require an e-mail reader that is UTF-8-capable.
- **Hybrid Analysis.** Unicode-mode applications now support Hybrid Analysis.

- **MDX.** MDX queries now support UTF-8-encoded character strings.
- **Data mining.** Unicode-mode applications now support data mining.
- **Query logging.** You can track queries on dimensions in databases in Unicode-mode applications. Create the query log settings file `dbname.cfg` in UTF-8 encoding and include the UTF-8 signature.

Triggers on Data Changes

Analytic Services Release 7.1 supports several new trigger features:

- On-update triggers and after-update triggers are now supported. The new after-update triggers fire only after all updates to the database are complete for a specific transaction.
- In MaxL, **create trigger** statements are supported as well as the existing **create or replace trigger** statements.
- For on-update triggers, you can now choose to add both old and new data values to the trigger spool file.
- You can set triggers for databases in Unicode-mode applications.

The new MaxL statements are as follows:

- **create after update trigger**
- **create or replace after update trigger**
- **create trigger**
- A new log value parameter on the **create on update trigger** and **create or replace on update trigger** statements provides support for adding both old and new data values to the spool file.

Note: On-update triggers are not supported for aggregate storage databases.

For additional information about this feature, see the *Essbase Analytic Services Database Administrator's Guide*. For specific details about the statements, see the *Technical Reference*.

Reduced Security-File Fragmentation

Fragmentation of the `essbase.sec` file can gradually develop when Analytic Services entities such as users, groups, applications, or databases are removed or changed. Too much fragmentation can slow down security-related performance. With this release, the security file is compacted each time the Agent is shut down. Additionally, you can compact the security file without shutting down the Agent, using any of the following methods:

- The COMPACT Agent command
- The `alter system compact security file;` MaxL statement
- The SECURITYFILECOMPACTIONPERCENT configuration setting

For additional information about this feature, see “Running Analytic Servers, Applications, and Databases” in the *Essbase Analytic Services Database Administrator’s Guide*. For information about the MaxL statement or the configuration-file setting, see the *Technical Reference*.

MaxL

MaxL includes the following enhancements:

- New grammar to support aggregate storage application and database operations
- Addition to the **display object** statement, for displaying only locked objects
- Additions to the **alter session** statement: one addition for setting MDX display options, and one addition enabling compression of the security file
- Additions to the **create trigger** statement, for specifying on-update triggers (the default) or after-update triggers

After-update triggers are supported only in traditional block-storage mode. For more information on the triggers feature, see [“Triggers on Data Changes” on page 5](#).

- Ability to view and delete data-monitoring logs (logs created by triggers) from within the MaxL Shell

For syntax and examples related to the MaxL changes, see the *Technical Reference*.

API

The Release 7.1 API supports the following new functions:

- Aggregate Storage Function:
EssCreateStorageTypedApplication()
- Trigger Functions:
 - EssDisplayTriggers()
 - EssMdxTrig()
 - EssListSpoolFiles()
 - EssDeleteSplFile()
 - EssDeleteAllSplFiles()
 - EssGetSpoolFile()
- C Outline API Function:
EssOtlQueryAttributesEx()

The Release 7.1 API provides performance improvements for EssOtlQueryMembers() and EssOtlQueryMembersByName(). The improvements enable API programs to request only the count of members rather than the entire member list. Alternately, the program can eliminate the member count calculation.

The following three C Outline API functions are no longer supported. Programs that contain these functions should be changed to avoid unnecessary errors:

- EssOtlGetNextAliasCombination()
- EssOtlDeleteAliasCombination()
- EssOtlAddAliasCombination()

Note: Release 7.1 introduces aggregate storage databases, which support multi-million member outlines. Compared to block storage databases, the number of members in aggregate storage databases may be several times larger. As a result, certain precautions are necessary when you migrate applications that use APIs that retrieve members from the server. For detailed information on the programming concerns for aggregate storage, see “Migrating to Release 7.1” in the *API Reference*.

Platform Support Changes

The following platform support was added for Release 7.0.1:

- Support for Windows 2003
- Support for Linux AS 2.1 for Intel

Documentation

The following topics discuss documentation improvements for Release 7.0 and Release 7.1:

- [“Documentation Installation Program” on page 30](#)
- [“Analytic Services Documentation Information Map” on page 30](#)

New Features in Release 7.0

The new features in Release 7.0 are described in the following topics:

- [“MetaData Security” on page 9](#)
- [“User Authentication Through the Hyperion Security Platform” on page 10](#)
- [“MaxL” on page 10](#)
- [“Report Writer” on page 15](#)
- [“Performance” on page 16](#)
- [“API” on page 17](#)
- [“Calculation” on page 18](#)
- [“Dimension Build” on page 19](#)
- [“Metadata Sampling” on page 20](#)
- [“Unicode Support” on page 21](#)
- [“Data Mining Support” on page 23](#)
- [“Monitoring Data Changes: Triggers” on page 25](#)
- [“Administration” on page 25](#)
- [“Delayed Free Space Recovery” on page 27](#)
- [“Server Threads Specified by Application” on page 28](#)
- [“Attribute Support for Hybrid Analysis Members” on page 29](#)
- [“Platform Support Changes” on page 29](#)
- [“Documentation” on page 30](#)

MetaData Security

Metadata security adds a layer of security for metadata (dimensions and members) in outlines, similar to read-only security for data cells. A new access level called *MetaRead* can be applied to dimensions or members. This level restricts access to dimensions or members during retrieval of an outline in a report.

Ways to apply metadata security include:

- Using the new access level, MetaRead, in the Filter Editor in Administration Services.

For more information, see the *Essbase Administration Services Online Help*.

- Using the `meta_read` keyword in the MaxL statements **create filter** and **alter filter**.

For more information, see the MaxL DDL section in the *Technical Reference*.

- Using a new API database access filter value called `ESS_ACCESS_METAREAD`.

For more information, refer to the *API Reference*.

User Authentication Through the Hyperion Security Platform

You can set up Analytic Server to use external authentication supplied by the Hyperion security platform. Although Analytic Services provided custom external authentication modules in earlier releases, the new security platform is common to all Hyperion products. Using the security platform to manage user accounts on Analytic Services provides these benefits:

- Analytic Services and Hyperion applications both can employ the existing corporate structure of user accounts, reducing administrative overhead.
- Single sign-on to Analytic Services from applicable Hyperion applications is now possible, eliminating the need for users to log on multiple times with multiple user names and passwords.

For more information, see the *Essbase Analytic Services Installation Guide* and the “Security Platform Reference” section of the *Technical Reference*.

MaxL

For a complete description of MaxL commands, refer to the *Technical Reference*. The following topics discuss new features and improved performance in MaxL:

- [“Filter Enhancements” on page 11](#)
- [“Improved STDOUT, STDIN, STDERR” on page 11](#)
- [“Improved Database-Related Grammar” on page 12](#)
- [“Display Location Aliases per Application” on page 12](#)

- “Display Port Information” on page 12
- “Full ESSCMD Functionality Available in MaxL” on page 12
- “Unicode-Related Operations” on page 14
- “Introduction of MDX” on page 14

Filter Enhancements

These are filter enhancements:

- Lines can be added to filter definitions using **alter filter**.
- A new level of filter access, `meta_read`, is added to **create filter** and **alter filter** statements.

Improved STDOUT, STDIN, STDERR

Most operating systems support three channels for input and output:

- STDIN (standard input channel)
- STDOUT (standard output channel)
- STDERR (standard error channel)

Most operating systems also provide command-line options for redirecting data generated by applications, depending on which of the above channels the data is piped through.

Prior to Release 7.0, the MaxL Shell (`essmsh`) printed all information to STDOUT unless the `spool` command was used. With Release 7.0, errors are now flagged as STDERR, allowing command-line redirection of errors using operating-system redirection handles.

You can also redirect STDERR or STDOUT independently to a MaxL output log, using the `spool` command. For syntax information, see documentation for the MaxL Shell `spool` command in the *Technical Reference*.

Improved Database-Related Grammar

These are the enhancements:

- The text of database notes is displayed as part of **display database**.
- Database free-space recovery is available using `alter database DBS-NAME recover freespace`

For more information about this feature, see [“Delayed Free Space Recovery” on page 27](#).

- Explicit database restructuring is available in **alter database**. See [“Explicit Restructure Command” on page 27](#).

Display Location Aliases per Application

The **display location alias** statement provides the added option of displaying location aliases associated with a particular application:

```
display location alias on application APP-NAME;
```

Display Port Information

The **display system** statement provides the added option of displaying agent port usage. This statement is the same as issuing the PORTS agent command:

```
display system ports;
```

Full ESSCMD Functionality Available in MaxL

The following table compares ESSCMD usage to MaxL usage, and the following conversions are now supported by the cmd2mxl utility. Only the comparisons pertaining to new MaxL grammar are listed. For the full mapping, see “ESSCMD Script Conversion” in the “MaxL DDL” section of the *Technical Reference*.

ESSCMD to MaxL Mapping

ESSCMD	Equivalent MaxL
GETDEFAULTCALC	query database DBS-NAME get default calculation
SETDEFAULTCALC	alter database DBS-NAME set default calculation as STRING;
LISTALIASES	query database DBS-NAME list alias_table
DISPLAYALIASES	query database DBS-NAME list alias_names in alias_table...

ESSCMD to MaxL Mapping (*Continued*)

ESSCMD	Equivalent MaxL
SETALIAS	alter database DBS-NAME set active alias_table...
LOADALIAS	alter database DBS-NAME load alias_table...
UNLOADALIAS	alter database DBS-NAME unload alias_table...
LISTOBJECTS	display object...
COPYOBJECT, RENAMEOBJECT, UNLOCKOBJECT	alter object OBJ-NAME...
DELETEOBJECT	drop object OBJ-NAME...
LISTLINKEDOBJECTS	query database DBS-NAME list lro
PURGELINKEDOBJECTS	alter database DBS-NAME delete lro
LISTLOCKS	display lock
REMOVELOCKS	drop lock
GETDBSTATS	query database DBS-NAME get dbstats...
GETATTRIBUTESPECS	query database DBS-NAME get attribute_spec
GETAPPINFO	display application
GETDBINFO	display database ... request_history
ESTIMATEFULLDBSIZE	query database DBS-NAME get estimated size...
GETCRRATE	query database DBS-NAME get currency_rate
SETHATRACE	alter database DBS-NAME set ha_trace to type...
VALIDATE	alter database DBS-NAME validate...
LISTFILES	query database DBS-NAME list file information
RESETPERFSTATS	alter database DBS-NAME set performance statistics...
GETPERFSTATS	query database DBS-NAME get performance statistics <type>
GETMBRCALC	query database DBS-NAME get member_calculation MBR-NAME
GETMBRINFO	query database DBS-NAME get member_info MBR-NAME

ESSCMD to MaxL Mapping (*Continued*)

ESSCMD	Equivalent MaxL
GETATTRINFO	query database DBS-NAME get attribute_info MBR-NAME
UPDATE	import database DBS-NAME data from data_string STRING...

Unicode-Related Operations

See “MaxL Support of Unicode” on page 22.

Introduction of MDX

Beginning with Release 7.0, MaxL has two functional domains:

- MaxL DDL is the data-definition language for Essbase Analytic Server.
- MDX is the data-manipulation language for Essbase Analytic Server.

Data Definition Language (MaxL DDL)

Data definition relates to structural control of a database system. This includes operations like creation, deletion, and updating of users, applications, databases, and database objects. Therefore, statements in MaxL DDL include verbs such as CREATE, ALTER, DROP, GRANT, and DISPLAY.

Data Manipulation Language (MDX)

Data manipulation relates to accessing the actual data within a database system. MDX in Analytic Services provides the ability to perform advanced data extraction and querying by means of statements that typically include the verb SELECT. The equivalent conceptual tool in Analytic Services would be Report Writer.

MDX is a joint specification of the XML for Analysis founding members. For more information about XML for Analysis, please visit: <http://www.xmla.org>

You can use either the MaxL Shell (essmsh) or MDX Editor in Administration Services to execute MDX queries. For example, the following query can be pasted into the MaxL Shell:

```
WITH MEMBER
    [Scenario].[Absolute Variance]
AS
    'Abs([Scenario].[Actual] - [Scenario].[Budget])'
SELECT
    { [Year].[Qtr1].children }
ON COLUMNS,
    { [Scenario].children, [Scenario].[Absolute Variance] }
ON ROWS
FROM
    Demo.Basic
WHERE
    ([Accounts].[Sales], [Product].[VCR],
    [Market].[San_Francisco])
```

When the above input is terminated with a semicolon and entered into the MaxL Shell, MaxL returns the results of the query in grid format. The semicolon is not part of the MDX language; it is an end-of-statement indicator to the MaxL Shell.

For more information about MDX, see the *Technical Reference*.

Report Writer

The following Report Writer commands are newly documented:

- GEN
- LEV
- <ACCOFF
- <ACCON

The GEN command selects all members from a specified dimension using the generation name. The LEV command selects all members from a specified dimension using the level name.

The <ACCOFF command turns off member accumulation when selecting members in a report script and the <ACCON command turns on member accumulation when selecting members.

For detailed information about these Report Writer commands, refer to the *Technical Reference*.

Performance

Hyperion provides performance improvements as discussed in the following topics:

- [“Improved Query Performance for Transparent Partitions” on page 16](#)
- [“Improved Performance for Large Outlines” on page 16](#)
- [“Improved Data Compression” on page 16](#)

Improved Query Performance for Transparent Partitions

Query performance has been improved for transparent partitions containing outlines that have a time dimension with Dynamic Time Series members. A new configuration setting, TARGETTIMESERIESOPT, turns on query optimization for transparent partitions.

For more information about TARGETTIMESERIESOPT, see the *Technical Reference*.

Improved Performance for Large Outlines

Data load times and overall index size have been improved for large outlines.

Improved Data Compression

In some instances, compression of the database has been improved. The most significant impact can occur on databases with large block sizes, where the blocks are highly sparse.

A new compression algorithm, Index Value Pair, has been added to Analytic Services. Index Value Pair addresses compression on databases with larger block sizes, where the blocks are highly sparse. You cannot assign Index Value Pair compression as the method to use. If either the RLE compression method or the bitmap compression method is assigned to the database, as Analytic Services writes blocks it analyzes the database and automatically uses Index Value Pair compression when appropriate.

The following table illustrates the available compression types the user can choose and the compression types that Analytic Services evaluates, then applies.

Chosen Compression Type	Evaluated Compression Type
None	None
Bitmap	Bitmap, Index Value Pair
RLE	RLE, Bitmap, Index Value Pair
zlib	zlib

API

Hyperion provides new features and enhancements for improved API functionality:

- A new API function that verifies the Rules file, `EssVerifyRulesFile()`
This function is especially useful when loading data through an external application.
- A new API function that can verify a formula during an outline edit session, `EssOtlVerifyFormula()`
This function corresponds to the Main C API function `EssVerifyFormula()`, which verifies formulas in a saved outline.
- A new API function that can verify an outline during an outline edit session, `EssOtlVerifyOutlineEx()`
This corresponds to the existing function, `EssOtlVerifyOutline()`, which verifies saved outlines.
- An upgrade of the Member Information structure to include additional information
Because this structure is retained locally during a client session, several time-consuming retrievals can be eliminated, saving the client session significant time.
- Many functions added to the Analytic Services Visual Basic API
These functions provide the same capabilities present in the Analytic Services C APIs
- New capability added to the Analytic Services API

This capability enables you to issue MDX statements from an application program. For more information about the capability, see [“MaxL” on page 10](#).

For a complete list of all the new Visual Basic API functions and the other new capabilities of the Analytic Services API, refer to the *API Reference*.

Calculation

These are the calculation enhancements:

- [“More Control of Currency Conversions after Loading Data” on page 18](#)
- [“Calculation Script Control over the Create Blocks on Equation Feature” on page 18](#)
- [“Calculation Script Control Over #MISSING Block Creation” on page 19](#)

More Control of Currency Conversions after Loading Data

Two new calculation script commands, SET CCTRACKCALC ON|OFF and CLEARCCTRACK, give you more control over the currency conversion process.

You can now update data in an individual currency partition and rerun currency conversions as required. You no longer need to use the RESETDB command and reload the entire database before rerunning a currency conversion.

SET CCTRACKCALC specifies whether Analytic Services uses the CCTRACK setting when calculating a currency conversion.

CLEARCCTRACK clears all conversion rate tables created by the CCTRACK setting. You can use this command with a FIX statement to clear the conversion rates for an individual currency partition.

For additional information about this feature, see [“Tracking Currency Conversions”](#) in the *Essbase Analytic Services Database Administrator’s Guide*. For specific details about the SET CCTRACKCALC and CLEARCCTRACK calculation commands, see the *Technical Reference*.

Calculation Script Control over the Create Blocks on Equation Feature

The new calculation script command, SET CREATEBLOCKONEQ ON|OFF, enables you to control whether blocks are created when non constant values are assigned to members of a sparse dimension. Whether or not blocks are created can affect database performance. This new calculation command temporarily overrides the database-level Create Blocks on Equation setting.

For additional information about this feature, see “Optimizing Calculations” in the *Essbase Analytic Services Database Administrator’s Guide*. For specific details about the SET CREATEBLOCKONEQ ON|OFF calculation command, see the *Technical Reference*.

Calculation Script Control Over #MISSING Block Creation

The new calculation script command, SET CREATENONMISSINGBLK ON|OFF, enables creation of new blocks during the execution of member formulas, including dense member formulas, any time the formula results in a value. The command ensures that only non empty blocks are stored after the calculation.

For additional information about this feature, see “Developing Calculation Scripts” in the *Essbase Analytic Services Database Administrator’s Guide*. For specific details about the SET CREATENONMISSINGBLK ON|OFF calculation command, see the *Technical Reference*.

Dimension Build

Hyperion provides certain rules file enhancements for building dimensions. These enhancements are discussed in the following topics:

- [“Changes to Existing Member Storage Properties” on page 19](#)
- [“Replacement of UDA Values” on page 19](#)
- [“Dimension Arrangement for Optimal Performance” on page 20](#)

Changes to Existing Member Storage Properties

A new S field member code enables you to change the storage property of members from the following previous storage properties: label only, Dynamic Calc, and Dynamic Calc and Store members to “store.”

Replacement of UDA Values

You can now use the UDA field type in the rules file to replace an existing UDA value. In prior releases, you could use the UDA field type to define new UDA values in the outline.

For additional information about use of the UDA field type in rules files, see “Setting Field Type Information” in the *Essbase Analytic Services Database Administrator’s Guide*.

Dimension Arrangement for Optimal Performance

You can now select an option in a dimension build rules file whereby Analytic Services automatically arranges the dimensions in hourglass order. In most situations, arranging the outline dimensions in hourglass order provides the most optimal dimension arrangement for calculation performance.

In hourglass order, the dense dimensions tagged as Accounts and Time are positioned at the top of the outline. The remaining dense dimensions are positioned next, from the largest down to the smallest. The sparse dimensions, which follow the dense dimensions, are arranged from smallest to largest. Attribute dimensions are always at the bottom of the outline and are not affected by this feature.

For additional information, read “Arranging Dimensions for Calculation Performance” in the *Essbase Administration Services Online Help*.

Metadata Sampling

Analytic Services is a multidimensional database engine that provides support for ad hoc analysis. Such analysis is entirely driven by the intuition of the analyst and can be time-consuming as data volume increases. Large cubes tend to have more dimensions and sometimes more levels, making hierarchical navigation very cumbersome.

Metadata sampling enables you to analyze on large cubes with a focus on data trends or approximate information in the initial stages. Because you query on a “sample” of your data, retrieval is quick. Metadata sampling enables you to zoom in on a portion of the vast amount of data in your database in a fraction of the time that it usually takes to analyze the whole database. You can view many samples in a small amount of time and make early decisions. Later, you can follow with organized data exploration.

With the introduction of Hybrid Analysis, you can store part of an Analytic Services cube in a relational database. Metadata sampling enables you to zoom in on all the data that you specify, whether it is in Analytic Services or in an underlying relational database. To support metadata sampling, Analytic Services adds these features:

- An API function: `EssGBeginSamplingZoomIn`
- A new menu command in the Essbase menu of the Spreadsheet Add-in for Excel: **Sample Data (Zoom In)**

This command acts as a toggle. A check mark next to this command indicates that sampling is enabled. When sampling is enabled, an approximate percentage of the Analytic Services database (or applicable relational data) is queried during Zoom In operations.

- A new text box on the Zoom tab of the Essbase Options dialog box in Spreadsheet Add-in for Excel: Sampling Percentage

Use this text box to indicate the approximate percentage of the Analytic Services database and relational database (if applicable) that you want to query during a Zoom In operation.

Spreadsheet Add-in for Excel retains all Zoom In settings when sampling is enabled.

For more information about the sampling API function, see the *API Reference*. For more information about the sampling feature in Spreadsheet Add-in, see the *Essbase Spreadsheet Add-in User's Guide for Excel* and Spreadsheet Add-in Help.

Unicode Support

Through its implementation of the Unicode standard, Analytic Services Release 7.0 enables employees of global businesses to view, in their own languages, company information stored in centralized databases.

Each Analytic Services application is designated as a non-Unicode-mode application (the default), or a Unicode-mode application. Analytic Services works with non-Unicode-mode applications the same way as it worked with applications in previous Analytic Services releases.

Administration Services, MaxL, and the API enable the migration of non-Unicode-mode applications to Unicode mode. Unicode-mode applications cannot be changed back to non-Unicode-mode. In Analytic Services Release 7.0, Unicode-mode applications do not support the following features:

- Hybrid Analysis
- Data mining
- Query logging
- SQL Interface
- MaxL Data Manipulation Language

This MaxL limitation means that triggers are not supported in Unicode-mode applications.

Character data in non-Unicode-mode applications is encoded to the code page common to server and clients. Character data in Unicode-mode applications is stored in UTF-8 encoding.

Unicode-mode applications also support additional characters for many objects such as application names and database names. See the “Limits” appendix in the *Essbase Analytic Services Database Administrator’s Guide* for details.

Unicode mode-applications accept both Unicode and non-Unicode-encoded input files. For Unicode-mode applications, the lengths of object names, such as names of applications, databases, members, and aliases, is in characters instead of bytes. In addition, the maximum lengths of application and database names have been increased.

To demonstrate Unicode-mode applications, the sample applications include the new, Unicode-mode application and database: Sample_U Basic. Member names in Sample_U Basic are in English.

Sample_U Basic includes four alias-table source files in different languages: nameschn.alt (Chinese), namesger.alt (German), namesjpn (Japanese), and namesrsn (Russian). To view member names in these languages in the Sample_U Basic outline, load these alias tables, then select the desired alias table before viewing the outline.

You can use the new Analytic Services Unicode Text Utility program, ESSUTF8, to convert standard text files to UTF-8 encoding, or to insert within a non-Unicode file, a header record that identifies the code page to which a non-Unicode file is encoded. For information about this utility, see the *Technical Reference*. By default, the agent log is in non-Unicode encoding. To enable viewing of application and database names that are in different character sets from the locale of the server, you can use the new configuration setting, UNICODEAGENTLOG, to set the agent log to UTF-8 encoding. For information about the UNICODEAGENTLOG configuration setting, see the *Technical Reference*.

Note: The Unicode implementation is transparent to users who do not need this feature.

For detailed information about working with the Analytic Services Unicode implementation, see the *Essbase Analytic Services Installation Guide*, the *Essbase Analytic Services Database Administrator’s Guide*, and the *Technical Reference*.

MaxL Support of Unicode

You can use MaxL DDL statements to perform the following Unicode-related tasks:

- Define the Unicode-related mode of Analytic Server using the `alter system` statement
- View the Unicode-related mode of Analytic Server using the `display system` statement
- Migrate non-Unicode-mode applications to Unicode mode using the `alter application` statement

- Create Unicode-mode applications using the `create application` statement
- Display the Unicode-related mode of an application using the `display application` statement

For details about these statements, see the *Technical Reference*.

Unicode Clients

Administration Services supports both non-Unicode-mode and Unicode-mode applications. See the *Essbase Administration Services Online Help* for additional information. You can use Application Manager only with non-Unicode-mode applications.

Note: Application Manager is not shipped with Release 7.0. You can access previous releases of Application Manager from the Hyperion Download Center using a pre-7.0 Analytic Services, such as Release 6.5.3 or higher.

You can use Essbase Spreadsheet Services to view data in both Unicode-mode applications and non-Unicode applications. Essbase Spreadsheet Add-in for Excel supports only non-Unicode-mode applications. You can use Spreadsheet Add-in for 1-2-3 from Analytic Services Release 6.5.3 or higher to view data in non-Unicode-mode applications.

New API functions have been added enabling C API support for both Unicode-mode and non-Unicode-mode applications. Existing custom-written client programs continue to work with non-Unicode-mode applications. Depending on how they are written, existing custom-written client programs need to be modified to work with Unicode-mode applications. For information about updating custom-written client programs written in C, see the *API Reference*.

Data Mining Support

A data mining framework has been added to this release of Analytic Services. Data mining is the process of searching through large amounts of data in an Analytic Services database to find hidden relationships and patterns. Data mining results can be descriptive, providing additional information about existing data; or predictive, forecasting future trends.

To support data mining, Analytic Services adds:

- A set of API functions
- A set of MaxL commands
- A data mining menu in Administration Services

The MaxL and Administration Services data mining interfaces enable users to:

- Add new Hyperion or third-party algorithms and modify or delete existing algorithms.
- Execute Hyperion or third party data mining algorithms.

Hyperion provides the following algorithms:

- Simple regression
Forecasts values based on the best straight line fit into the data. The result helps detect trends in data.
- Multilinear regression
Captures dependencies between a value of interest and several contributing factors. The result can be used to model the potential behavior of the subject when the factors are varied.
- Clustering
Arranges items of interest into groups that appear to form similar patterns. The result can lead to better understanding of market segmentation, provide insight into possible market factors, and serve as a basis for other predictive tools.
- Association rules
Provides market-basket analysis with results in the form “*nn%* of the users who buy *Product X* also buy *Product Y*”.
- Decision tree
Determines simple rules for making decisions, such as whether to accept or deny a loan application. The algorithm constructs a collection of yes-or-no questions. The next question to ask depends on the answer to the current question. The result looks like a tree and the goal is to arrive at a final yes-or-no answer.
- Neural net
Constructs a multi-layered set of nodes that model certain numerical dependencies between input and output data. The nodes within one layer connect to the next layer and thus make up a net that resembles human neurons.

For detailed information regarding data mining, read the “Mining an Analytic Services Database” chapter of the *Essbase Analytic Services Database Administrator’s Guide*, the “Mining an Analytic Services Database” topic in the *Essbase Administration Services Online Help*, and the MaxL commands section of the *Technical Reference*.

Monitoring Data Changes: Triggers

The Analytic Services triggers feature enables you to monitor data changes in a database efficiently. If data breaks the rules that you have specified, Analytic Services can send an e-mail (to a user or system administrator) or log the information in a file. For example, you might want to send an e-mail to the Sales manager if sales in the Western region fall below sales for the equivalent month in the previous year.

You can create and delete triggers using these MaxL statements:

- **create or replace trigger**
- **alter trigger**
- **display trigger**
- **drop trigger**

For information on creating and removing triggers, read the “Monitoring Data, Applications, and Databases” chapter of the *Essbase Analytic Services Database Administrator’s Guide*, and the MaxL documentation in the *Technical Reference*.

Administration

Hyperion provides additional or improved administration functionality. The enhancements are discussed in the following topics:

- [“Application Log Enhancements” on page 26](#)
- [“Multiple Agent Enhancements” on page 26](#)
- [“Optimize Outline for Batch Calculation” on page 26](#)
- [“Explicit Restructure Command” on page 27](#)

Application Log Enhancements

Application log messages have been enhanced with the name of the object used to execute an operation, if the object resides on an instance of Analytic Server:

- For calculations, the name of the calculation script used to perform the calculation is added.
- For data loads, the rules file and the data file name are added.

Note: Analytic Services supplies the load rule name only if the client making the request (for example, Administration Services) is at the same release level as Analytic Services; that is, Release 7.0.

- For a report, the report script name is added.

Multiple Agent Enhancements

The process for a client to connect to multiple agents on a single machine has been greatly simplified. Previously, to connect to a different agent, a client had to change the port number in its configuration file, stop, then start again. This same process had to be repeated to go back to the original agent.

Now a client can dynamically connect to an agent by specifying the port number of the agent at the time of connection.

For more information regarding connecting to multiple agents, see the chapter, “Running Analytic Servers, Applications, and Databases” in *Essbase Analytic Services Database Administrator’s Guide*.

Optimize Outline for Batch Calculation

As you design, develop, and tune a database outline, you can use the new Administration Services Optimize Outline feature to apply standard design principles to the outline for optimizing performance of batch calculations. Be aware that the standards applied using this feature may not optimize your outline because of factors unique to your deployment such as calculation scripts, end-user workloads, and hardware environment.

The Administration Services Optimize Outline feature rearranges dimensions. It also changes the storage property of members to Label Only or Dynamic Calc, in the certain circumstances. The following members are changed to Dynamic Calc:

- In sparse dimensions, generation 1 members that have eight or fewer children
- All parent members in dense dimensions
- Members tagged as Two-Pass

Parents, grandparents, and so on, of Two-Pass members in dense dimensions are tagged as Two-Pass.

The following parent members are changed to Label Only:

- Parent members with the first child aggregating up to them
- Parent members with all children tagged for no consolidation

For additional information, see “Optimizing Outlines for Batch Calculation” in the *Essbase Administration Services Online Help*.

Explicit Restructure Command

The ability to restructure the database explicitly (to eliminate or reduce fragmentation) has been added. Previously, the only way to restructure the database was by adding a member (and then deleting it) or to force a structure change in a block.

The means to restructure the database include these enhancements:

- New grammar in the **alter database** MaxL statement:

```
alter database DBS-NAME force restructure
```
- An API function, `EssOtlRestructure()`.

Delayed Free Space Recovery

Improvements have been made to the free space recovery process. Free space recovery is the most expensive part of database recovery. Database recovery takes place any time you load an application that ends abnormally.

Changes have been made to the free space recovery process, as follows:

- Performance has been improved significantly.
- Free space recovery, which happened automatically before, is now delayed by default. The user must explicitly trigger free space recovery. A configuration option has been added to specify whether or not to delay free space recovery by default.

Note: All database functions run as usual whether you recover free space or not. The advantage of recovering free space is that you can reuse disk space marked as free in the data files. The disadvantage is that free space recovery is time-consuming.

- Explicit free space recovery has been added with a new MaxL grammar for **alter database**, as follows:

```
alter database DBS-NAME recover freespace
```

- The GETDBSTATS command adds two lines of information about free space recovery, as follows:

```
Free Space is Recoverable           : Yes/No
Estimated Bytes of Recoverable Free Space : nnn
```

- If free space is recoverable, the block counters are estimates and do not necessarily match the number of existing blocks.
- A new C API function, EssRecoverDbFreeSpace, has been added.
- The following ESB_DBSTATS_T structure has been changed:
 - The NonExclusiveLockCount and ExclusiveLockCount fields are obsolete and have been removed.
 - The AverageFragmentationQuotient and BytesOfRecoverableFreeSpace fields have been added. AverageFragmentationQuotient is the percentage of space within the data file that is free space, or that is not used by Analytic Services. BytesOfRecoverableFreeSpace is set to -1 if free space recovery has not been delayed and does not need to be run at this time. Otherwise, it is a nonnegative number that is an estimate of how many bytes of free space can be recovered by running RecoverDbFreeSpace.
 - The NonMissingLeafBlocks and NonMissingNonLeafBlocks fields provide estimates only of the number of missing blocks.

For more information regarding free space recovery, see the free space recovery topic in the chapter, “Ensuring Data Integrity” in *Essbase Analytic Services Database Administrator’s Guide*. For more information about the **alter database** syntax, see the MaxL section of the *Technical Reference*. For information about the DELAYEDRECOVERY setting, see the “Configuration Settings” section of the *Technical Reference*.

Server Threads Specified by Application

Previously, the SERVERTHREADS configuration setting could only specify a single server thread value for all applications, meaning each application on the server had to use the same value. Now, users can choose to specify a different number of server threads for certain applications.

For more information about SERVERTHREADS syntax, see the *Technical Reference*.

Attribute Support for Hybrid Analysis Members

Hybrid Analysis enables you to integrate a relational database with an Analytic Services database. This integration leverages the scalability and size of the relational database with the conceptual and reporting capabilities of the multidimensional database. In this release, the power of Hybrid Analysis has been amplified by enabling you to perform data operations and analyses on attribute dimensions. The attribute dimension members can also be associated with members that are in a relational source.

For more information, read “Accessing Relational Data with Hybrid Analysis,” in the *Essbase Analytic Services Database Administrator’s Guide*.

Platform Support Changes

Hyperion supports the following platforms for the server:

- AIX 5L Version 5.1 and 5.2
- Sun Solaris 8 and 9 (Sun OS 5.8 and 5.9)
- HP-UX Version 11.0 and 11i
- Windows 2000
- Windows XP (Personal Essbase only)

Hyperion supports the following platforms for the client on Windows 2000 and XP and Citrix Metaframe XP with Windows 2000 server:

- Excel 97, Excel 2000, and Excel XP
- Lotus 1-2-3 Release 9.6 using Analytic Services Release 6.5.3 or higher Spreadsheet Add-in for Lotus 1-2-3

Note: Spreadsheet Add-in for Lotus 1-2-3 is not shipped with Analytic Services Release 7.0. To access previous releases of Spreadsheet Add-in for Lotus 1-2-3, go to the Hyperion Download Center at <https://hyperion.subscribenet.com>.

Documentation

The following topics discuss documentation improvements for Release 7.0:

- [“Documentation Installation Program” on page 30](#)
- [“Analytic Services Documentation Information Map” on page 30](#)

Documentation Installation Program

Starting with Release 7.0, Analytic Services documentation is installed separately from the software product installation. The documentation installation program is available from the Download Center or from CD in the Analytic Services package. For more information on the documentation installation process, see the *Essbase Analytic Services Installation Guide*.

Analytic Services Documentation Information Map

You can access all Analytic Services documentation from the *Essbase Analytic Services Information Map*, an HTML file that is installed with the documentation. Access the information map in either of these ways:

- On Windows, select Start > Programs > Hyperion Solutions > Essbase Analytic Services > Documentation
- Navigate to `$ARBORPATH/docs/esb_infomap.htm`.

Note: Procedures for performing administrative tasks are documented in *Essbase Administration Services Online Help*.

Where to Get More Information

Except for the most recent information, each topic in this booklet is described in more detail in the documentation.

For answers to questions about the product, contact your authorized technical support provider or

Hyperion Solutions Customer Support
Phone: 203-703-3600 (outside the U.S.A.)
Phone: 877-901-4975 (in the U.S.A.)
Internet: <http://support.hyperion.com>

Visit the Hyperion Solutions Web Site Home Page at <http://www.hyperion.com>.

Copyright 1998–2004 Hyperion Solutions Corporation.

All rights reserved.

U.S. Patent Number: 5,359,724

Hyperion, Essbase, and the Hyperion “H” logo are registered trademarks, and Hyperion Solutions is a trademark of Hyperion Solutions Corporation.

All other brand and product names are trademarks or registered trademarks of their respective holders.

No portion of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser’s personal use, without the express written permission of Hyperion Solutions Corporation.

Notice: The information contained in this document is subject to change without notice. Hyperion Solutions Corporation shall not be liable for errors contained herein or consequential damages in connection with the furnishing, performance, or use of this material.

Hyperion Solutions Corporation
1344 Crossman Avenue
Sunnyvale, CA 94089

Printed in the U.S.A.