



## **General Concepts**

---

Last updated: 28/10/2016

# Table of Contents

About This Document.....	4
1 Architecture Overview.....	5
1.1 Architecture Components.....	5
1.1.1 Lightstreamer Server.....	6
1.1.2 Lightstreamer Adapters.....	6
1.1.3 Lightstreamer Clients.....	7
1.2 Summary of the Components.....	8
2 High-level Concepts.....	9
2.1 Multichannel Transport.....	9
2.2 Message Routing.....	10
2.3 Optimized Delivery.....	12
2.3.1 Dynamic Throttling.....	12
2.3.2 Other Optimizations.....	13
2.4 Scalability.....	14
2.5 Security.....	14
2.6 Monitoring.....	15
3 Main Concepts.....	16
3.1 Glossary.....	16
3.2 Data Model and Subscription Modes.....	18
3.2.1 MERGE Mode.....	19
3.2.2 DISTINCT Mode.....	21
3.2.3 COMMAND Mode.....	22
3.2.4 RAW Mode.....	23
3.2.5 Compatibility of Modes.....	23
3.3 Bandwidth and Frequency Management.....	24
3.3.1 Bandwidth Management.....	24
3.3.2 Frequency Management.....	24
3.3.3 Applying the Limits.....	24
3.3.4 Filtered and Unfiltered Modes.....	25
3.3.5 The ItemEventBuffer.....	26
3.3.6 ItemEventBuffer in RAW Mode.....	27
3.3.7 ItemEventBuffer in MERGE Mode.....	27
3.3.8 ItemEventBuffer in DISTINCT Mode.....	27
3.3.9 ItemEventBuffer in COMMAND Mode.....	28
3.3.10 Resume of Settings for the ItemEventBuffer.....	28
3.3.11 Considerations on the Usage of the Buffer.....	29
3.3.12 Selectors.....	29
3.4 The Preprocessor.....	30

3.4.1 Snapshot.....	30
3.4.2 Prefilter.....	31
3.4.3 Preprocessor and Modes.....	31
3.5 Roles in the Choice of Parameters.....	32
4 The Adapters.....	33
4.1 The Metadata Adapter.....	33
4.1.1 Authentication and Authorization Examples.....	35
4.2 The Data Adapter.....	36
4.2.1 An Internal Data Adapter: The Monitoring Data Adapter.....	39
5 Mobile Push Notifications.....	44
5.1 MPN Subscriptions.....	45
5.1.1 The Formatting Syntax.....	46
5.1.2 The Trigger Expression.....	46
5.1.3 MPN Subscription ID and Coalescence.....	47
5.2 The Database Provider.....	48
5.2.1 The Database in a Clustered Configuration.....	49
5.2.2 DBMS Compatibility.....	49
5.3 The MPN Provider.....	49
5.3.1 APNs.....	50
5.3.2 GCM.....	50

## About This Document

---

This document provides a comprehensive picture about the general architecture of Lightstreamer, by describing its main features and design principles.

The document is intended for Architects, Developers, System Integrators and System administrators who have the responsibility of design, develop, integrate, deploy and administer every Lightstreamer based solution.

This document is structured in the following sections:

- Architecture Overview outlines the building blocks of the Lightstreamer Architecture.
- High-level Concepts Introduces some of the characteristics of the Lightstreamer technology.
- Main Concepts provides fundamental notions needed to understand the internal mechanisms of Lightstreamer.
- The Adapters illustrates how custom Adapters work.
- Mobile Push Notifications describes role and functioning of the special module designed to integrate both Apple Push Notification Service and Google Cloud Messaging.

# 1 Architecture Overview

Lightstreamer is a real-time engine that delivers real-time data over HTTP and WebSockets to/from any types of clients (HTML pages, native mobile applications, desktop applications, etc.), as well as native push notifications to mobile clients.

By implementing unique features such as **Bandwidth & Frequency Throttling** and **Adaptive Streaming**, Lightstreamer handles the distribution of live data to the clients in a reliable, light, and efficient way.

Lightstreamer is able to pass through any firewall and proxy by using many sophisticated mechanisms without the risk of security policy blockages.

This chapter provides an overview of the Lightstreamer architecture.

*NOTE: Some of the features described in this document may be available on some Lightstreamer editions only. Please refer to the product matrix to know the details (<http://www.lightstreamer.com/products>).*

## 1.1 Architecture Components

The Lightstreamer Server is the core component of the Lightstreamer technology. It is a stand-alone process that runs in a Java Virtual Machine. It handles the connections with the clients and dispatches real-time data back and forth between the clients and any back-end system.

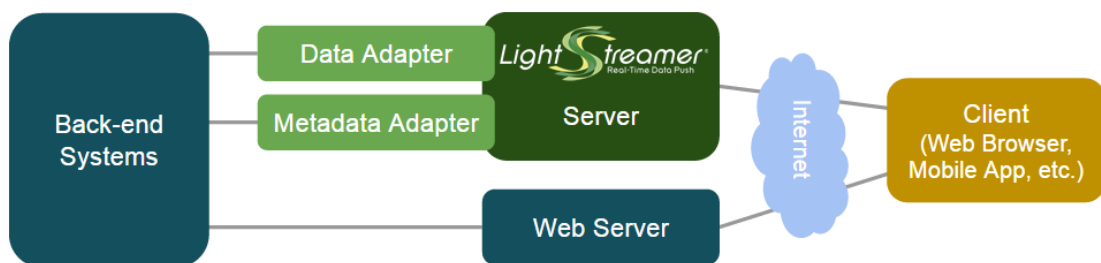


Figure 1- Lightstreamer General Architecture.

Although Lightstreamer comes with an embedded Web server to serve static resources, which makes the product immediately useful for demo purposes, the deployment architecture shown above is normally employed in real production scenarios.

Typically, the Lightstreamer Server and a common Web Server (Apache, nginx, IIS or any other web server) are deployed in the DMZ network, accessible over the Internet by clients. Lightstreamer Server gets the data and metadata from the back-end systems that are protected by the second firewall.

The clients (web browser, mobile applications, desktop applications, etc.) get synchronous data from the Web server (*pulled data*, such as web pages or any request/response interaction) and asynchronous real-time data from the Lightstreamer Server (*pushed data*). The clients can also send real-time data to the Lightstreamer Server. This way, bi-directional streaming (or two-way push) is fully supported.

In addition, native push notifications (Apple APNs and Google GCM) are supported for mobile clients (see §5).

It is possible to create a **cluster** of Lightstreamer Servers by using any Web load-balancing appliance, in order to achieve **load-balancing** and **fail-over**. If one of the machines in the cluster fails, the clients will reconnect to another machine. The **recovery** is transparent to the users, while full data coherency is maintained. See the "Clustering" documentation for more details

### 1.1.1 Lightstreamer Server

---

The Kernel of Lightstreamer Server is the "off-the-shelf" part of the system. It takes direct control over the operating system's TCP/IP stack and handles the real-time flow of data for each client to optimize the data transmission as much as possible. The Kernel implements several important features to offer an efficient and reliable push service, such as:

- **Scalability.** The architecture of the Kernel is based on a staged event-driven architecture, that allows many thousands of concurrent connections to be sustained by a single server CPU.
- **Firewall and Proxy Friendly.** The use of HTTP and HTTPS over standard ports to deliver data to the Clients allows the traffic to pass through the strictest firewalls, proxies and routers. Some proxy servers exist that block any form of streaming, because they try download the full content of a document before forwarding it to the client. In this cases, the **Stream-Sense** features of Lightstreamer automatically detects the issue and switches to a high-efficiency smart polling mode.
- **Bandwidth Control.** For each user a maximum bandwidth can be allocated, dedicated to the streaming channel. For example, if it is required that a certain user cannot exceed a bandwidth of 10 kbps, the Kernel will filter the data in such a way as to ensure that the streaming connection for that user always remains below 10 kbps.
- **Frequency Control.** A maximum update frequency can be allocated for each user/item combination. For example, it is possible to configure the profile of a certain user so that they do not receive more than 2 updates per second for a certain subscribed item (piece of data).
- **Adaptive Streaming and Congestion Control.** The Kernel automatically detects situations of congestion on the network, heuristically slowing down or suspending the dispatch of data until the connection is again fully serviceable. Data aging is avoided.
- **TCP-Level Optimization.** The Kernel has direct control over the composition of the TCP packets. Instead of delegating the aggregation of data in packets to the operating system (using the *Nagle* algorithm), the Kernel itself decides on each occasion the optimum composition of each TCP packet, with the objective/trade-off of reducing the latency before data is dispatched as well as minimizing the number of dispatched packets.
- **Prefiltering, Selector Mechanism, Event Customization.** Advanced features to exploit the data flow before it is delivered to the clients.
- **JMX Manageability** Any management console compliant with Java Management Extensions can communicate with the JMX agent included in Lightstreamer Kernel to collect thorough data about the Server status and to manage the system.

### 1.1.2 Lightstreamer Adapters

---

Lightstreamer Adapters are custom server-side components attached to the Lightstreamer Server, whose role is to interface the Kernel with any data source, as well as to implement custom authentication and authorization logic. An *Adapter Set* is made up of a Metadata Adapter and one or multiple Data Adapters, both developed by implementing provided interfaces.

Adapter APIs are currently provided for **Java**, **.NET**, **Node.js**, and **Python**. In the case of Java, Adapters can run either in-process with the Lightstreamer Server or in an external process (on the same machine or on a different machine). .NET and Node.js Adapters always run in external processes. In addition, Adapters can be developed in any language to talk to the Lightstreamer Server via plain **TCP sockets** through the **Adapter Remoting Infrastructure** (see the specific documentation for full specification of the protocol).

### Lightstreamer Metadata Adapter

The role of a Metadata Adapter is to manage the policies applied to a user's session. It controls the life cycle of user sessions, performing the following tasks:

- validation of authentication credentials sent by the client
- assigning of the granted quality of service to each user, in terms of allocated bandwidth and update frequency
- interpretation, validation and authorization of subscription requests, made up of items and field names

### Lightstreamer Data Adapter

The Kernel of Lightstreamer is integrated with any data feed (the source of real-time information) by means of a Data Adapter.

Each Data Adapter is a connector that interfaces Lightstreamer Server with a data source. It receives a flow of data from the back-end systems (information providers, data feeds, databases, application servers, CRMs or any other legacy platforms) and injects it into the Lightstreamer Kernel for controlled delivery to individual users.

The system integrator that implements the Data Adapter is free to use any technology to integrate the Kernel with the data feed. However, it is preferable to use middleware that provides asynchronous paradigms, such as message-oriented systems, so as not to break the asynchronous chain that goes from the feed to the user's client. In any case, it is also possible to use polling techniques to get updates from the data feed (e.g. polling a file or a database).

Because Lightstreamer is able to manage an arbitrary number of different Data Adapters and Metadata Adapters (grouped together into Adapter Sets), *it is possible to integrate heterogeneous sources of information while maintaining a single point of access to the streaming/push channel.*

## 1.1.3 Lightstreamer Clients

---

Lightstreamer client libraries are provided for many different platforms, which expose APIs to integrate any application with the Lightstreamer Server. The available client APIs include:

- Web (compatible with any browser, including older browsers and mobile browsers; supports frameworks like AngularJS and React, as well as hybrid frameworks, such as PhoneGap and Electron)
- Android
- Apple (iOS, macOS, tvOS)
- Microsoft (.NET, Excel, Silverlight, Windows Phone, WinRT)
- Java SE
- Node.js (for both server-side code and React Native apps)
- Python
- Unity
- Adobe (Flash, Flex, AIR)

- Java ME and BlackBerry
- Generic client (open protocol)

## 1.2 Summary of the Components

To resume, Lightstreamer comprises the following main components:

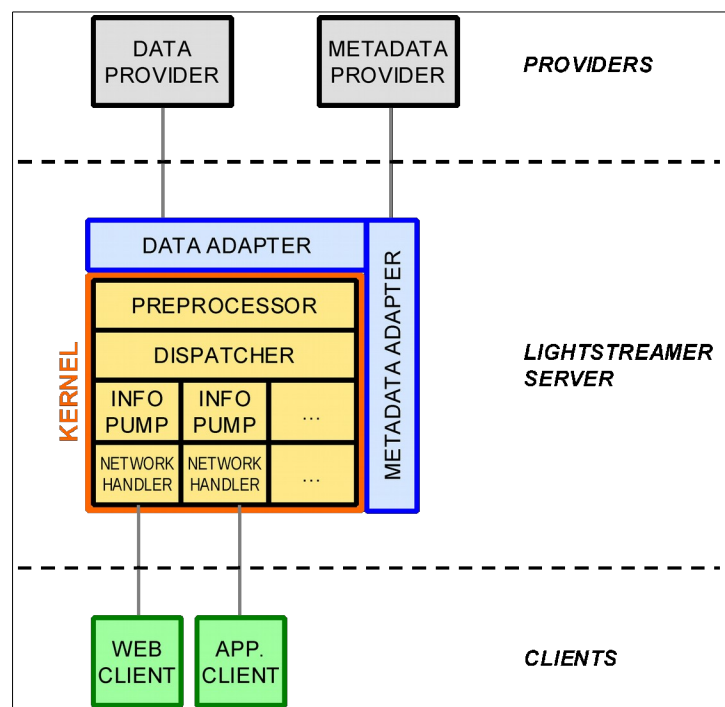


Figure 2- The Components of Lightstreamer Server.

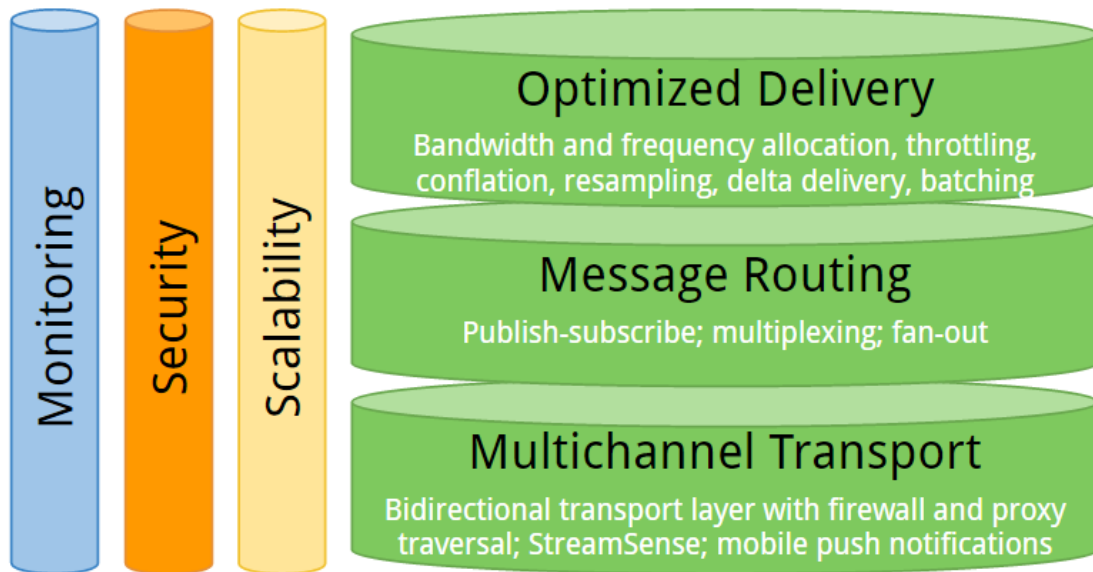
- **Data Provider** – Data feed that provides Lightstreamer with real-time data updates.
- **Metadata Provider** – Metadata feed that provides Lightstreamer with authentication information, permissioning policies and description of data (for example, items and fields).
- **Lightstreamer Data Adapter** – Custom component that integrates Lightstreamer Kernel with the Data Provider.
- **Lightstreamer Metadata Adapter** – Custom component that integrates Lightstreamer Kernel with the Metadata Provider.
- **Lightstreamer Kernel** – The core of the server, which in turns is comprised of the following sub-components:
  - **Preprocessor** – performs Prefiltering and pre-elaboration on update events.
  - **Dispatcher** – forwards real-time updates from the Data Adapter.
  - **InfoPump** – data structure allocated by the Kernel for each activated session.
- **Lightstreamer Client** – Any kind of application (Web-based, mobile, tablet, desktop, as well as server processes) that receives/sends real-time data from/to Lightstreamer Server through a TCP/IP network (using only HTTP and/or WebSockets). A Client subscribes to a set of data that the Server publishes.



## 2 High-level Concepts

From a functional perspective, Lightstreamer is made up of three logical layers: Multichannel Transport, Message Routing, and Optimized Delivery.


Scalability, Security, and Monitoring are three essential properties.



### 2.1 Multichannel Transport

The transport layer offers the upper layer the abstraction of a reliable network transport that works in any situation. Web protocols (HTTP and WebSockets) are used. This, together with several sophisticated mechanisms, enables bi-directional real-time communication through any kind of proxies, firewalls, and other network intermediaries.

The **Stream-Sense** mechanism implements a fast detection of the best transport on a per-client basis, with a sequence of automatic fallbacks. For example, the JavaScript client library behaves like this:



Whatever is the transport automatically chosen for the communication with each client, the same exact set of features is guaranteed. In particular, a high-performance bidirectional channel is provided in all the cases, with in-order guaranteed message delivery and automatic batching from client to server. In other words, Lightstreamer enriches HTTP when sending messages from the client to the server:

- Messages are acknowledged explicitly
- Lost messages are retransmitted automatically
- Out-of-order messages are reordered automatically
- Underlying socket is kept open for reuse via reverse heartbeats
- Multiple requests are automatically batched, to highly reduce the number of HTTP round trips


In addition to Web transports, Lightstreamer supports native transports for **push notifications** to mobile clients, where the protocol and infrastructure provided by Apple (**APNs**) and Google (**GCM**) are leveraged (see §5).

## 2.2 Message Routing

The Multichannel Transport layer offers reliable message delivery on the Internet. But this is not enough to build flexible and complex applications. Some message routing facility is needed to provide a convenient way to deliver the right data to the right recipients. For this reason, on the top of the Web Transport layer, Lightstreamer offers a powerful Message Routing layer.


Lightstreamer provides a **publish and subscribe** (pub-sub) paradigm, to decouple message producers from message consumers.

A client subscribes to items (the basic piece of information) while specifying a schema (a set of fields for each item):




(The "S" logo represents the Lightstreamer Server in this set of figures)

When an item is subscribed to by a client for the first time, the subscription is propagated to the Data Adapter, which should start publishing updates for that item. Any subsequent client subscribing to the same item does not cause any new request to the Data Adapter. When the last client unsubscribes from that item, the Data Adapter is notified it can stop publishing updates for that item:




Thanks to **on-demand publishing**, no data needs to be produced when no actual recipient exists, saving on systems resources.

The Lightstreamer Server receives from the Data Adapter the initial snapshot for each item (for example, the current value for all its fields), followed by the real-time updates. It then sends the data to each client, based on its subscriptions, in a multiplexed fashion, on the top of a single connection:



Thanks to the flexibility of the subscription-based mechanism, any message routing scenario is actually supported. From **unicasting** (where a message targets only one client), to **broadcasting** (where a message reaches all the clients), passing through **multicasting** (where a message reaches a group of clients).



For example, the top part of the figure above shows a case where the same item is subscribed to by one million clients. The Data Adapter needs to publish updates to that item only once, while the **fan-out** is done by the Lightstreamer Server.

The bottom part of the same figure shows the opposite case. An item is subscribed to by one client only, so that every client gets its own **personal message flow**.

Broadcast, multicast, and unicast items can be freely mixed as part of the same client session.

Lightstreamer implements an **asymmetric pub-sub** paradigm:



This means that the publishers are on the server side (connected to the Data Adapters) and the subscribers are on the client side. This highly optimizes the cases where there is some massive publishing from typical data feeds. In many scenarios, the data feed is completely different from the data consumer (due to topology, to protocol, or to the business model).

That being said, clients can still publish data, by sending direct messages to the Data Adapter (intermediated by the Metadata Adapter for security reasons):



This way, the Data Adapter has always the responsibility of publishing the actual updates for each item, irrespective of the fact that such updates are coming from a server-side data feeds or from the clients. Furthermore, in some cases the Data Adapter will want to validate, reformat, and process the data before publishing it. This is the extreme flexibility provided by the asymmetric pub-sub model, compared to traditional symmetric pub-sub systems, where publisher and subscribers are all peers (and some complex ESB might be needed to validate and adapt the data flows).

## 2.3 Optimized Delivery

Given a reliable and scalable Multichannel Transport abstraction and a flexible and secure Message Routing facility, a higher-level layer can be implemented, to optimize as much as possible the data delivery. The Delivery Optimization layer of Lightstreamer provides a very powerful set of features to save on bandwidth, reduce latency, and cope with the Internet unpredictability.

### 2.3.1 Dynamic Throttling

Based on the nature of the data, series of updates to an item can be filtered, to reduce their frequency (message rate), via:

- **Queueing** (buffer all the messages and deliver them at the proper rate)
- **Resampling** (discard some of the messages)
- **Conflation** (discard some of the messages while preserving the latest value for each field).

Imagine you have a sensor that measures the current temperature 100 times per second and you want to deliver the measurements in real time to a software application. If you just want to display the value to a human user, perhaps 10 updates per second are more than enough. Or, even better, you might want to deliver as many updates as possible without saturating the link bandwidth. This is where resampling and conflation are extremely useful.

For each subscription of each client, Lightstreamer allows to define how data can be filtered, with several parameters. Filtering is then applied on the fly to the data stream based on a number of static and dynamic variables.

There are three main variables that Lightstreamer takes into consideration to decide how to throttle the data flow:

#### 1) Bandwidth allocation

For each client, a maximum bandwidth can be allocated to its multiplexed streaming connection. Lightstreamer guarantees that such max bandwidth is never exceeded, whatever is the original data update rate. Max bandwidth can even be changed dynamically during the life of a client session.

## 2) Frequency allocation

For each subscription of each client, a maximum update frequency (message rate) can be allocated. Lightstreamer guarantees that such max frequency is never exceeded, whatever is the original data update rate. Max frequency can even be changed dynamically during the life of a client session.

## 3) Real bandwidth detection

Internet congestions are automatically detected by the Lightstreamer Server and the actual available bandwidth at any point in time for each client is estimated.

In addition to these, there are other variables taken into considerations, either directly or indirectly (CPU cycle availability, batching configuration, etc.).

- A live demo is available, which shows Bandwidth and Frequency allocation in action: <http://demos.lightstreamer.com/BandwidthDemo/>

In a nutshell, Lightstreamer is able to continuously adapt the data flow to the network conditions. With data that can be resampled, Lightstreamer can send a different stream to each client, where the number of samples depends on the available bandwidth.

Suppose a mobile client connected over a tiny-bandwidth network with terrible packet loss and a desktop client connected with a broad-bandwidth connection subscribe to the same item. This item might contain the 3D coordinates of a user in a multiplayer game, or the telemetry data of an airplane, or the prices of a financial instrument. In all the cases, each client will receive an amount of data compatible with its network capacity. The first user above will probably receive less updates per second than the second user. But both will receive up-to-date data, not aged data that stacked up in a queue. Resampling done on the fly for each individual user is one of the most powerful features of Lightstreamer.

### 2.3.2 Other Optimizations

---

Several other techniques are employed by Lightstreamer to optimize the data delivery.

#### **Data is aggregated efficiently within TCP packets.**

The TCP Nagle's algorithm is disabled (TCP\_NODELAY option) and Lightstreamer uses its own algorithms to decide how to pack data into TCP segments. A trade-off between latency and efficiency can be configured. What is the maximum latency you can accept to forge bigger packets and decrease the overheads (both network overheads and CPU overheads)? You can answer this question for each application and game and configure it in Lightstreamer via the *max\_delay\_millis* parameter. The higher the acceptable delay, the more data can be stuffed into the same TCP segment (**batching**), increasing overall performance. For extreme cases, you can use 0 for *max\_delay\_millis*.

#### **Delta delivery.**

With delta delivery, for subsequent updates to an item, only the actually changed fields (delta) are sent. For example, if a user subscribes to 20 different fields for each stock (last, bid, ask, time, etc.), only a few of them really change at every market tick. Lightstreamer automatically extracts the delta and the provided client-side library is able to rebuild the full state. Upon initial connection, Lightstreamer sends a full snapshot for all the subscribed items.

### Lightweight protocol.

Lightstreamer avoids using JSON, XML, or any other verbose protocol, which carries large amounts of redundant data with each event (for example, the field names), resulting in increased bandwidth usage. Lightstreamer uses a position-based protocol, which reduces the overhead to a minimum. The actual payload accounts for most of the bandwidth.

A description of these and other techniques, with some examples, is available in a Lightstreamer article, which describes a gaming use case but has general validity: <http://blog.lightstreamer.com/2013/10/optimizing-multiplayer-3d-game.html>

## 2.4 Scalability

---

The internal architecture of Lightstreamer Server makes it extremely scalable. It is based on a concurrent staged event-driven architecture with non-blocking I/O.

Performance and scalability depend on many different variables, but as a very rough figure, Lightstreamer has been tested on a single box with: one million connections with low frequency traffic; tens of thousands connections with very high frequency traffic. In other words, the actual scalability depends most of all on the message throughput.

Lightstreamer provides graceful degradation of the quality of service. If the server CPU saturates, due to wrong capacity planning or to an unexpected spike, the Lightstreamer Server will not block. It will simply start to reduce the update rate evenly on all the connections, while avoiding data ageing at the same time (thanks to adaptive throttling; see below).

The Lightstreamer Server can scale both vertically and horizontally.

**Vertical scalability:** The Lightstreamer Server automatically leverages all the CPUs and cores that are available on that machine, by spawning the right number of threads and avoiding any critical global synchronization point.

**Horizontal scalability:** Multiple Lightstreamer Servers can be clustered by using any standard Web Load Balancer. See the "Clustering" documentation for more details.

## 2.5 Security

---

Lightstreamer offers pluggable security policies via the Metadata Adapter. Basically, every authentication and authorization decision can be delegated to the Metadata Adapter, in a very fine-grained fashion. For example, upon the client session creation, the authentication credentials received from the client, together with further information (http headers, network properties, etc.) are passed to the Metadata Adapter for validation. The Metadata Adapter is free to query whatever back-end system it wants to decide whether to authenticate a user or not.

Similarly, each subscription done by a client, as well as any request for some quality of service (QoS), such as granted bandwidth or update frequency, is authorized through the Metadata Adapter.

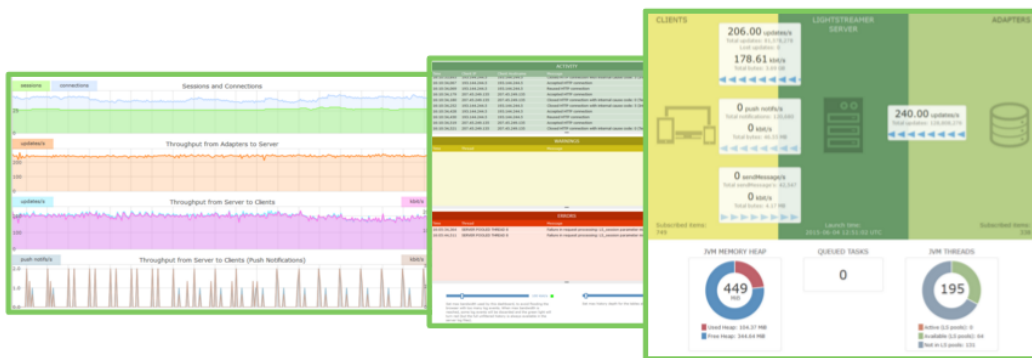
See §4.1 for more details.

## 2.6 Monitoring

Lightstreamer offers several means to monitor the server.

A world-class **monitoring dashboard** is available in Lightstreamer 6.0. Based on HTML, it is configured to work out of the box and provides many real-time metrics and statistics on the behavior of the Lightstreamer Server. The dashboard is organized in four tabs:

- The first tab shows a visual representation of the Lightstreamer architecture, with several overlays containing the metrics (connections, sessions, throughput, memory, thread pools, etc.).
- The second tab displays real-time graphs of the main metrics.
- The third tab contains the logs.
- The fourth tab hosts the "JMX Tree", where the monitoring objects can be browsed and inspected.



The Monitoring Dashboard is a convenient means of watching the server for both developers and systems administrators. By default, it is available at the "/dashboard/" URI of the Lightstreamer Server.

A very fine-grained configurable **logging** facility is available, with several categories, levels, and appenders. It is based on logback (<http://logback.qos.ch/>).

Extensive metrics are exposed via a **JMX interface**, to hook into any third-party application management tool. See the JMX Extension SDK for full details.

### 3 Main Concepts

#### 3.1 Glossary

The picture below shows a graphical representation of the hierarchy between a Session and an Item. Definitions of each term are given below the figure.

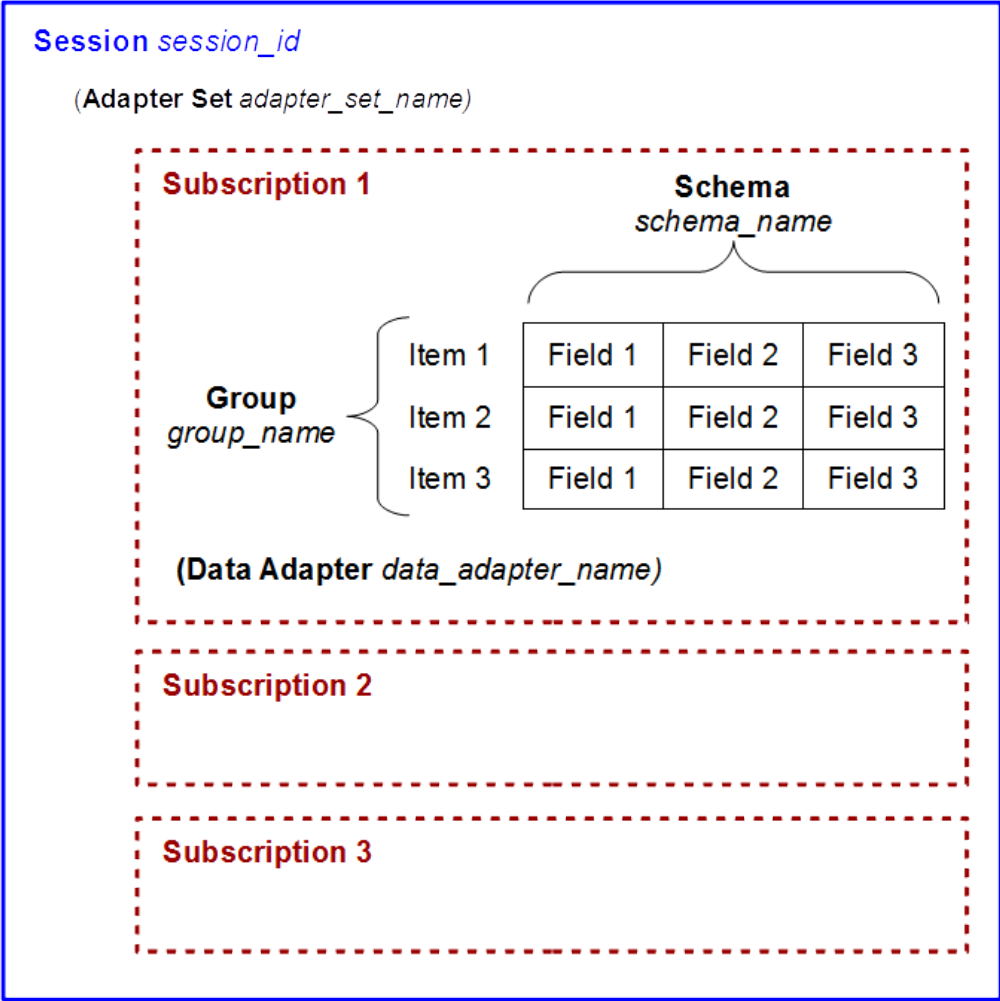


Figure 3- The logical hierarchy between a Session and an Item.

#### Adapter Set

It is made up of one Metadata Adapter and one or multiple Data Adapters. Multiple Adapter Sets can be plugged onto Lightstreamer Server.



## Session

An entity created in Lightstreamer Server targeted to serve a specific Client. Each session is identified by a unique **sessionId**. Each session is associated to a single **Adapter Set**.

## User

A user to whom a session is bound. Each session is bound to one and only one user.

## Item

Information entity whose state is defined by a set of **fields** whose values can change over time. In other messaging systems, items are often referred to as "subjects", "topics", or "channels".

Each item is identified by a unique **item name** (whose syntax depends on the supplier Data Adapter). There exist both "general items", which can be subscribed by any user, and "personal items", each of which can be subscribed only by a specific user or by a group of users. This enables unicast messaging and multicast messaging, as opposed to broadcast messaging. The distinction between these categories of items is enforced by the Metadata Adapter, which can validate each subscription based on the user identity.

Examples:

- item name = "MSFT.NASDAQ" (a full quote of Microsoft stock on NASDAQ exchange)
- item name = "user123-prv-msg" (a private user message).

## Field

A string pair comprised of a **field name** and a **field value**.

Examples:

- field name = "bid"; field value = "3.05" (the bid element of a stock quote)
- field name = "ask"; field value = "3.04" (the ask element of a stock quote)
- field name = "message"; field value = "Hello World!" (a message in a chat room)

## Snapshot

The current full state of an **item**. For example, in MERGE mode, the snapshot is the set of **fields** that define the complete state of an item at that time.

## ItemEvent

A set of **fields** sent from a Data Adapter to Lightstreamer Kernel to update the state of an **item** at a certain time.

### ItemUpdate

A set of **fields** sent from Lightstreamer Kernel to a Lightstreamer Client to update the state of an **item** at a certain time.

### UpdateEvent

A set of **ItemUpdates** sent from Lightstreamer Kernel to a Lightstreamer Client to update the states of all the subscribed items.

### Field Schema

A set of **field names** that have been subscribed to by a Client in relation to an **item**. A Field Schema is identified either by an array of field names or by a unique string that should be known to the Client and to the Metadata Adapter.

### Item Group

A set of **items** subscribed to by a Client with a common **field schema** and a common supplier **Data Adapter**. An Item Group is identified either by an array of item names or by a unique name that should be known to the Client and to the Metadata Adapter.

### Subscription

An association among an **Item Group**, a **Field Schema**, and a **Mode**. It may also be referred to as a **Table** in the context of some client libraries.

### Mode

The way the **itemEvents** of an **item** are handled in Lightstreamer Kernel for a Client. It can be: "**MERGE**", "**DISTINCT**", "**COMMAND**", and "**RAW**". The mode is indicated by the Client during subscription.

### Selector

A predicate to be applied to the **itemEvents** in order to filter them. It can be associated to a **Subscription**, so that it will be applied for all **items** in the subscribed **Item Group**. The selector is identified by a symbolic name that should be known to the Client and to the Metadata Adapter, which is wholly responsible for the filtering.

## 3.2 Data Model and Subscription Modes

---

The role of a Data Adapter is to map the original *data model* of a Data Provider onto the Lightstreamer data model. As anticipated, the Lightstreamer data model is based on items, which are made up of fields. Items can be interpreted in different ways, depending on the **subscription mode**.

With **MERGE** subscription mode, an item represents a row in a table. Real-time updates to that item are used to update the contents of the cells (fields) for that row.

With **DISTINCT** subscription mode, an item represents a list of events. Real-time updates to that item are used to add lines to that list (where each line is made up of fields).

With **COMMAND** subscription mode, an item represents a full table. Real-time updates to that item are used to change the contents of that table, by adding rows, removing rows, and updating cells.

By combining these subscription modes, it is possible to map any kind of data. It is always advised to try to leverage the Lightstreamer fields as much as possible, so that the field-level delivery optimizations can take place. But there is total freedom on what items and fields should represent. So, if tabular data representation is not suitable and some complex data structure should be delivered, it is possible to embed it into a single field.

In addition to these modes, there exists **RAW** subscription mode, where no assumption is made on the nature of the data, leading to less optimization in data delivery.

The following sections delve into the details of each subscription mode.

### 3.2.1 MERGE Mode

---

Many data sources can be defined as filterable. **Data filtering** is a possibility offered by the intrinsic nature of some types of information. All information sources that over time generate updated versions of the same data, more or less frequently, can quite easily be subjected to filtering algorithms.

The following are examples of filterable data feeds:

- **Stock prices.** The price of a stock ("last price") varies continuously during the course of a stock market session. Once the market has generated a new price, the previous one automatically becomes obsolete. So, in the case where prices are generated with high frequency, one could decide not to communicate each and every variation. It is obviously indispensable to maintain the consistency of the data presented in the application, despite the filtering. For example, if on the same page there is the price of a covered warrant that usually gets updated once every hour, together with the price of a stock listed on the NASDAQ, which generates 30 updates a second, it is obviously unthinkable to lose the one update of the first in order to communicate more updates of the second.
- **Measurements (monitoring, metering, telemetry, etc.).** A probe of any sort (e.g. a physical probe to measure temperature or a software probe for network management) produces a whole series of measurements which are samples of the quantity that the probe is designed to monitor. These samples can be subsequently re-sampled to reduce their frequency even further.

Lightstreamer manages the filterable data feeds by means of the **MERGE mode**.

In this case, the source provides updates to a persisting state of an item. The absence of a field in an itemEvent is interpreted as "unchanged data". Any "holes" must be filled by copying the latest available value for each field.

It is important to note that filtering is not mandatory, even in MERGE mode (see §3.3.4).

#### **Conflation (Merging)**

Not all the itemEvents from the Data Adapter need to be sent to the Client (depending on throttling configuration). With conflation, if a previous itemEvent is discarded, the changes brought by it are still included in the newer itemEvents.

Let's imagine that each update is a small sheet divided into cells, with the current field values written in each cell. By referring to the stock price example, the fields are "last", "last size", "bid", and "ask". The sheets fall down one after another from the data source and stack on the ground, as shown in the following picture:




Figure 4 - The updates falling down and stacking on the ground.

Each update carries only the fields whose value has changed. The other cells are transparent. The first sheet (Update #1) carries all the field values. The second one (Update #2) carries only "ask". The third one (Update #3) carries only "last" and "last size".

Upon the first update, the current item state is given as follows:

*Item = [last=10.55, last size=125, bid=10.54, ask=10.58]*

When the second sheet stacks upon the first one, a new sheet is created, that carries "last", "last size", and "bid" from the first sheet and "ask" from the second sheet. This leads to an updated current item, whose state is given by merging old (black) and new (red) values:

*Item = [last=10.55, last size=125, bid=10.54, ask=10.57]*

This means that if for any reason (bandwidth constraint, frequency allocation, adaptive streaming, etc.) the first two updates cannot be sent to the client, only the resulting *merged* update ("sheet") will be sent.




Figure 5 – The first and second updates are merged.

When the third sheet stacks upon the first two, a resulting sheet carrying "last" and "last size" from the third sheet, "ask" from the second sheet, and the "bid" from the first sheet is created. The final item state is:

*Item = [last=10.54, last size= 500, bid=10.54, ask=10.57]*

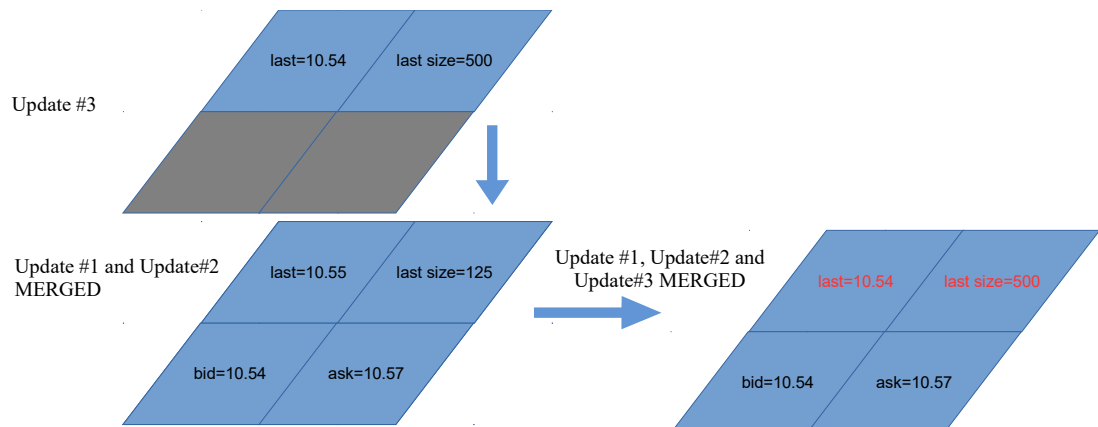


Figure 6– All the three updates are merged together.

In this case, the server has the possibility of sending one update only to the client instead of three. But as mentioned before, such filtering only occurs when actually necessary, based on several combined conditions that determine throttling.

### 3.2.2 DISTINCT Mode

There are data sources that do not allow conflation. In these cases, the feed does not produce data that replaces the previous one, but each event is distinct from the previous one.

The following are examples of distinct data feeds:

- **News headlines.** Press agencies generate fresh news items throughout the day. Typically, a real-time news visualization system shows a series of headlines in chronological order (from the latest to the oldest) which runs on receipt of each new headline. So when a fresh item of news is received, it does not eliminate the previous one; it gets listed alongside it.

- **Chat messages.** Messages in a chat room should follow one each other without being “conflated” or resampled.

Lightstreamer manages distinct data feeds by means of the **DISTINCT mode**.

In this case, the source provides events of the same type. The itemEvents coming from the Data Adapter must be sent to the client unchanged. DISTINCT data can be either unfilterable or filterable (see §3.3.4).

### 3.2.3 COMMAND Mode

There exist many application scenarios where it is required not only to push changes to an item content in terms of updating one or more fields, but even changes to the list of items itself, in terms of adding and removing rows. This is the so called “**meta-push**”.

The following are examples of meta-push feeds:

- **Stock portfolio.** The value of stocks in a portfolio keeps changing based on the market, so some fields keep being updated. But, at the same time, selling or buying some stocks results in removing or adding rows to the portfolio table.
- **Chat participants.** The list of users in a given chat room keeps changing, due to users entering and exiting the room. The state of users in the room can change as well, for example from “active” to “idle”.

Lightstreamer manages the meta-push feeds by means of **COMMAND mode**

In COMMAND mode, usually a single item is used for the whole list or table (e.g. the whole portfolio). The itemEvents are interpreted as **commands** that indicate how to progressively modify a list. In the schema there are two fields that are required to interpret the itemEvent.

1. The “**command**” field contains the command associated with the itemEvent, having a value of “ADD”, “UPDATE” or “DELETE”.
2. The “**key**” field contains the key that unequivocally identifies a line of the list generated from the Item.

The “key” and “command” fields must never be missing from an itemEvent.

In the following figure, a graphical representation of an Item handled in COMMAND Mode is shown:

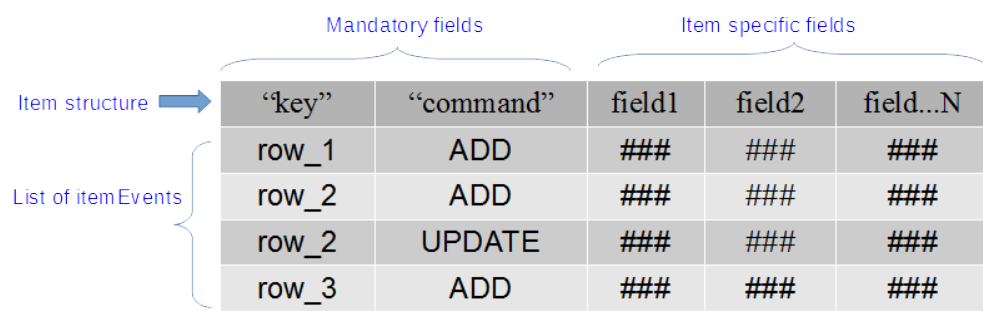


Figure 7 - Item in COMMAND Mode.

Not all the itemEvents coming from the Data Adapter need to be sent to the client (because filtering is possible). Successive events regarding the same “key” (but not necessarily received consecutively from the Data Adapter) can be merged according to the following rules:

- an UPDATE merges with previous ADD or UPDATE;
- a DELETE removes the previous ADD or UPDATE (when removing ADD, the DELETE itself is removed; when removing UPDATE, the DELETE remains);

- an ADD removes the previous DELETE and becomes an UPDATE.

For the itemEvents with commands of "ADD" or "DELETE" the "holes" (in the other fields) are propagated to the Client. For the itemEvents with commands of "UPDATE", the "holes" are filled by copying the values of the field corresponding to the last message of ADD or UPDATE with the same key.

### Two-level Push

To achieve a better separation of concerns, the COMMAND mode enables the so called "Two-level" push behavior too. In this way, it is possible to decouple the ADD/DELETE logic from the UPDATE logic, by means of two different kinds of items:

- the first-level item, based on COMMAND mode, which will manage the "first-level" subscriptions in terms of changes to the list of keys;
- the second-level items, which will manage the "second-level" subscriptions in terms of updates to the values for each single key.

Each time the Data Adapter pushes an "ADD" command for the first-level item with a new key to add a row, the client uses this information to automatically subscribe to an **underlying second-level item** in MERGE mode. Similarly, when a "DELETE" command is received to remove a row from the list, the client automatically unsubscribes from the underlying item. This mechanism is natively supported by some client libraries.

This way, the first-level item represents the full table (determines which rows are part of the table at every moment); for each row, there is a specific second-level item that provides the field values for such row.

The first- and second-level items may be supplied by different Data Adapters. For example, in the stock portfolio mentioned above, it is easier to have one Data Adapter managing the changes in the portfolio contents, and another Data Adapter managing the updates to the prices coming from the market.

### 3.2.4 RAW Mode

---

In RAW Mode, all the itemEvents coming from the Data Adapter must be sent to the client unchanged. The only constraint that applies is the bandwidth allocation. RAW mode acts as a "dumb pipe", where no aggregation of events is done.

### 3.2.5 Compatibility of Modes

---

In principle, the same item can be subscribed to in more modes by the same client or by different clients. However, each item can only be handled in one among the MERGE, DISTINCT, and COMMAND modes, depending on the kind of source, which restricts the possibilities somewhat. On the other hand, every item can be handled in RAW mode. This gives rise to 4 possible logical families of items, as follows:

1. RAW
2. RAW | MERGE
3. RAW | DISTINCT
4. RAW | COMMAND

The expected modes for each Item have to be specified by the Metadata Adapter, so that client requests for other modes can be rejected. This is done by implementing the *modeMaybeAllowed* method, which must obey the above restrictions. Failing to do so will cause the Server to handle each item according to the mode specified in the first subscription request it receives for it, and to just ignore subsequent requests for conflicting modes.

The Metadata Adapter can further restrict the allowed modes for each Item, either overall or on a per-client basis. For instance, when either MERGE or DISTINCT or COMMAND Mode is supported, it could disallow RAW mode in order to enforce filtering.

## 3.3 Bandwidth and Frequency Management

---

### 3.3.1 Bandwidth Management

---

Each stream connection is characterized by a maximum bandwidth (**maxBandwidth**), which is an absolute and binding limit. The bandwidth limit is ensured even across consecutive polling connections.

The maxBandwidth depends on the user and the client and is derived as follows:

$$\text{maxBandwidth} = \min(\text{allowedMaxBandwidth}, \text{requestedMaxBandwidth})$$

where:

- **allowedMaxBandwidth** is the maximum bandwidth the Metadata Adapter allocates for a specific user's session.
- **requestedMaxBandwidth** is the maximum bandwidth the client can request at the opening of the stream connection.

It is possible to omit these restraints (or set them as "unlimited").

### 3.3.2 Frequency Management

---

Each item subscribed to by a client is characterized by a maximum frequency (**maxFrequency**) to send updated events to that client. The frequency limits are ensured even across consecutive polling connections.

The maxFrequency is determined as follows:

$$\text{maxFrequency} = \min(\text{allowedMaxFrequency}, \text{requestedMaxFrequency}, \text{grantedMaxFrequency})$$

where:

- **allowedMaxFrequency** is the maximum frequency the Metadata Adapter allocates to the item/user pair.
- **requestedMaxFrequency** is the maximum frequency the client can indicate at the moment of the subscription of the Item Group containing that item.
- **grantedMaxFrequency** is an upper limit on the frequency that may be imposed as a limitation in some specific editions of the product.

It is possible to omit these restraints (or set them as "unlimited").

### 3.3.3 Applying the Limits

---

Maximum bandwidth and maximum frequency constraints act on different levels: while the bandwidth constraint is applied to the whole session, the frequency constraint is applied to each item individually. Both constraints set an upper bound, which is dynamically managed by the Lightstreamer Server.



A Client can specify the requestedMaxBandwidth and/or the requestedMaxFrequency limits in order to protect itself from update rates too high to be sustainable. A client may also need to ensure that the itemUpdates for an item don't exceed a maximum frequency for application-level reasons (think to updates that are meant to be displayed in a cell). By setting a requestedMaxFrequency, the client will have the Server apply the limit before sending the updates, hence saving bandwidth. Moreover, a client may also specify the "unlimited" values to mean that the maximum actual parameters can be decided on the Server side.

On the other hand, the Metadata Adapter can specify limits (i.e. allowedMaxBandwidth and allowedMaxFrequency) in order to protect the Server from peaks of activity. While doing this, it has the opportunity of granting different service levels to different clients based on the involved user.

Finally, the license terms may impose limits (i.e. grantedMaxFrequency) that are only meant to force a reduced quality of service on limited versions of Lightstreamer.

When bandwidth and frequency limits are set, some ItemEvents originated by the Data Adapter may have to be *filtered*. This is done for each Item independently, and in different ways depending on the subscription mode. However, even with no limits set, filtering can occur, at Server discretion, in case of resource limitations or for uninformative updates (again, depending on the mode).

#### 3.3.4 Filtered and Unfiltered Modes

---

With respect to the maximum frequency, the management Modes can further be specialized in "filtered" and "unfiltered" modes:

- The **unfiltered mode** is activated on the client side by requesting **unfiltered dispatching** at the moment of the subscription. This is accomplished by specifying the value "unfiltered" for the requestedMaxFrequency parameter. Such mode prevents filtering at all for an item and obviously implies requesting no frequency limits. However, the request is accepted only if the allowedMaxFrequency allocated by the Metadata Adapter for an item/user pair is "unlimited"; otherwise the user cannot subscribe the with unfiltered dispatching.
- The **filtered mode** is activated on the client side by specifying for the requestedMaxFrequency parameter a number or the value "unlimited" at the moment of the subscription.

The management of requestedMaxFrequency and allowedMaxFrequency is provided only for the following subscription modes: **MERGE**, **DISTINCT**, **COMMAND** (in **COMMAND** the frequency limits apply to the UPDATE commands sent for each key). For **RAW** mode, unfiltered dispatching (and so unfiltered mode) is always implied.

On the other hand, grantedMaxFrequency applies to all items, even if subject to unfiltered dispatching.

The following table shows a final resume of the allowed combined settings to determine the maxFrequency parameter:

Filtered/Unfiltered	requestedMaxFrequency	allowedMaxFrequency	Output maxFrequency
Filtered mode	Number	Number	min(requestedMaxFrequency, allowedMaxFrequency)
	"unlimited"	Number	allowedMaxFrequency
	Number	"unlimited"	requestedMaxFrequency
	"unlimited"	"unlimited"	Unlimited
Unfiltered mode	"unfiltered"	"unlimited"	Unlimited

Table 1 - Allowed combinations for maxFrequency parameter.

- An **online example** with full source code, showing dynamic changes to bandwidth and frequency is available at:  
<https://github.com/Lightstreamer/lightstreamer-example-BandwidthandFrequency-client-javascript>

### 3.3.5 The ItemEventBuffer

For each item subscribed to in a session, an **ItemEventBuffer** (which is a FIFO buffer) is created in the InfoPump instance which is dedicated to that session (an InfoPump is a data structure allocated by the Lightstreamer Kernel for each session). More precisely, distinct itemEventBuffers are set up for each single subscription: for instance, if the same item is subscribed twice in the same session, two buffers will be prepared.

The itemEventBuffer contains the itemEvents extracted from the supplier Data Adapter that are still to be forwarded to the Client, because of bandwidth and/or frequency constraints. Only the requested fields are retained, according to the field schema indicated during subscription.

Each element will be sent to the Client as part of an updateEvent (except for the case when the buffer is full, more details further on).

An itemEventBuffer is characterized by the following parameters

- **mode:** mode of interpreting the new data in arrival (**RAW, MERGE, DISTINCT, COMMAND**).
- **dispatching switch:** "filtered" or "unfiltered", as described in paragraph 3.3.4
- **bufferSize:** maximum buffer dimension (whose value can also be "unlimited"); it can be specified for MERGE and DISTINCT modes, unless unfiltered dispatching is in place. It refers to the number of events waiting for dispatching and is relative to each single subscription.

The choice of mode is affected by the nature of the data and is indicated during item subscription. It impacts on the way events can be packed into updateEvents and on the way buffer size limits apply.

The unfiltered mode presupposes that all events participate in the evolution of the client state and no event can be suppressed. In this case, more events can be released in a single updateEvent. However, grantedMaxFrequency and bandwidth limits can still lead to the need for a buffer, in case of event bursts. For this reason, when in unfiltered mode, an **unlimited** bufferSize is always granted.

In all the cases where the bufferSize is unlimited, the Kernel is free at runtime to leverage a “robustness mechanism” which puts a limit on the maximum dimension of the buffer. This means that the Kernel can block a further growth of buffer size to protect itself, by discarding events. When this happens for items in RAW mode or in unfiltered dispatching, the Kernel must also signal to the Client that one or more events were lost; it is up to the Client to decide whether to abort the subscription to that item, or to maintain it with the knowledge that not all events are being received.

In general, when the buffer is full and a new itemEvent arrives from the supplier Data Adapter, this will still enter the buffer, forcing the rejection of the element at the top of the buffer (the first element that entered, i.e. the oldest).

### 3.3.6 ItemEventBuffer in RAW Mode

---

In RAW Mode, more than one itemUpdate can be release in a single updateEvent, depending on bandwidth limits. The bufferSize is always **unlimited**. Any data loss is signaled to the Client.

### 3.3.7 ItemEventBuffer in MERGE Mode

---

In **MERGE filtered mode**, only one itemUpdate event can be released in a single updateEvent (unless the allowed maxFrequency is very high or “unlimited”). This way, as a matter of fact, the frequency of releasing of the updateEvents (which depends on various factors, including a maxBandwidth limit and the network roundtrip time) acts as a further frequency constraint and may enforce further filtering.

A buffer dimension greater than 1 allows it to absorb event bursts without overlaying them, in cases where it is preferred to receive as many updates as possible without filtering. Exceeding the buffer size limits leads to the silent suppression of some events, applying the merging algorithm. The default size is 1.

In **MERGE unfiltered mode**, more than one itemUpdate can be released in a single updateEvent. In this case the buffer size must be **unlimited**. Any data loss is signaled to the Client.

### 3.3.8 ItemEventBuffer in DISTINCT Mode

---

In **DISTINCT filtered mode**, only one itemUpdate event can be released in a single updateEvent (unless the allowed maxFrequency is very high or “unlimited”).

A buffer with a size greater than 1 allows event bursts to be absorbed, in order to send them to the Client in such a way that it is compatible with the assigned maximum frequency and without loss. Exceeding the buffer size limits leads to the silent suppression of some events. The default size is unlimited.

In **DISTINCT unfiltered mode**, more than one itemUpdate can be released in a single updateEvent. In this case the buffer size must be **unlimited**. Any data loss is signaled to the Client.

### 3.3.9 ItemEventBuffer in COMMAND Mode

---

In **filtered COMMAND mode**, for each "key", the events belonging to the "key" are released in a way similar to the MERGE mode with a buffer size of 1. Exceeding this 1-sized buffer leads to event "matching" and merging. Hence, as said, in COMMAND mode frequency limits apply to the various keys independently. If the "robustness mechanism" must be leveraged, it will try to filter out events of "UPDATE" type; if any events of "ADD" or "DELETE" type must be discarded, then this is also signaled to the client, as some problems may occur, particularly during snapshot management (see below). However, note that as long as the number of active keys is less than the protection limit, the buffer cannot grow beyond the limit, regardless of the event rate.

In **unfiltered COMMAND mode**, each single event is granted not to be suppressed and is delivered to the client as is, unless the buffer has grown over the protection limit; in such case, any data loss for "ADD", "UPDATE" and "DELETE" events is notified to the Client.

On the other hand, the unfiltered COMMAND mode does not ensure that events pertaining to different keys are sent to the clients in the same order in which they are received from the Data Adapter (this may be particularly evident when grantedMaxFrequency limits, which apply to the various keys independently, are involved). In fact, usually, one should expect that updates for different keys are not related to one another.

#### Ordering

There are cases in which strict ordering could be useful; one case is handling ordered lists (classifications, charts, results, etc.). If the sort criterion were based on a field representing a "score", then we would only need to sort the column of the table on the client side. If, however, the sort criterion were more complex, then the client could be sent a field that contains the current position in the list. But the move of a single element could imply the need for an UPDATE for each line of the table and this could be very heavy and redundant. To optimize the handling, we can assume that the client will be sent only the UPDATE for the element that is moved (or that enters or exits), leaving the client the task of modifying the logical and physical position of the other elements. This would not work properly if update events could be filtered out or even flipped across different keys.

At the current stage, the further restriction of preserving update order across keys for unfiltered COMMAND mode can only be activated on a Server-level scope through a proper configuration flag. However, in that case, any grantedMaxFrequency limit would apply to the item as a whole; in fact, in cases like the one above, the benefits of using COMMAND mode instead of RAW mode would be minimal.

### 3.3.10 Resume of Settings for the ItemEventBuffer

---

The table below summarizes the possible combined settings as detailed in the previous paragraphs:

Mode	Buffer Size	ItemUpdate per UpdateEvent relative to the same item
RAW	Unlimited	1..N (but 1 with grantedMaxFrequency)
MERGE filtered	1..unlimited (default =1)	1 (but 1..N if frequency is unlimited)
DISTINCT filtered	1..unlimited (default = unlimited)	1 (but 1..N if frequency is unlimited)
COMMAND filtered	Unlimited	1..N with different keys
MERGE unfiltered	Unlimited	1..N (but 1 with grantedMaxFrequency)
DISTINCT unfiltered	Unlimited	1..N (but 1 with grantedMaxFrequency)
COMMAND unfiltered	Unlimited	1..N (but only with different keys with grantedMaxFrequency)

**Table 2 - Settings for the ItemEventBuffer.**

As said, if a grantedMaxFrequency limit is set because of the use of a limited edition, it applies to all itemEventBuffers, regardless that unfiltered dispatching is in place. Moreover this also extends to all itemEventBuffers the limitation of 1 ItemUpdate per UpdateEvent relative to the same item (or to the same "key", for unfiltered COMMAND).

Consider that the presence of a tight maxBandwidth or grantedMaxFrequency limit poses further and possibly strong restrictions to the ItemUpdate throughput, even on items in RAW or unfiltered mode.

In particular, for RAW or unfiltered items, the internal buffer may grow and since the bufferSize is unlimited, the ItemUpdates may be released with significant delays and the "robustness mechanism" may eventually take place. It is an application responsibility to ensure that the update rate from the supplier Data Adapter cannot exceed the outbound rate.

### 3.3.11 Considerations on the Usage of the Buffer

In normal cases, the Server should not need to use the buffer, as it should be able to send all events in real time. When subscribing to an item in filtered mode, the subscription settings should also request a small buffer size.

On the other hand, if a subscription is made in a mode that does not allow filtering, then the event rate for the item should not be higher than the expected dispatch rate. However, if the incoming event rate for an item is steadily higher than the dispatch rate, then the buffer usage grows with negative impacts both on dispatching delays and on Server heap memory usage.

### 3.3.12 Selectors

A predicate on the itemEvents (called a Selector) can be associated to the ItemEventBuffer. It is applied before itemEvents enter the ItemEventBuffer and can prevent some itemEvents to enter the buffer.

Typically, selectors can be used to extract subsets of events, based on the "key" value in COMMAND mode or on some special attribute in RAW and DISTINCT modes. Selectors may also be used to suppress unimportant updates in all modes. Selector can be applied on a per-user basis.

NOTE: In COMMAND mode, events suppression should be consistent with the key lifecycle, ensuring that ADDs and DELETes on the same key are interlaced (always starting with an ADD, if the snapshot has been requested) and that UPDATES follow ADDs for the same key, but not DELETes.

NOTE: In MERGE mode, the suppression of a snapshot event may cause the client to receive the first event for the Item only some time after the subscription, despite the snapshot information has been requested.

## 3.4 The Preprocessor

---

The Preprocessor has the task of filling the empty fields in the events, maintaining snapshots and performing Prefiltering.

### 3.4.1 Snapshot

---

When a client subscribes to an item, it may request the **initial snapshot**, that is, a way to bootstrap the current item state without having to wait for the asynchronous real-time updates.

Only for the items handled in MERGE, DISTINCT, and COMMAND modes, the Preprocessor can maintain snapshots and make them available to the InfoPump when it requests them (when the Client requests a snapshot to the InfoPump). RAW mode does not support the snapshot.

**MERGE:** The snapshot is the current value for each field. Consisting of an itemEvent which is progressively modified or extended by itemEvents coming from the Data Adapter.

**DISTINCT:** The snapshot is the recent history of updates to the item. Consisting of a list with a predefined maximum length, which contains the latest itemEvents received from the supplier Data Adapter. All the events belonging to the snapshot are sent to the client in an *unfiltered dispatching* fashion, despite the frequency and buffer size limits operating for normal updates.

**COMMAND:** The snapshot is the shortest sequence of commands needed to build the current state of the table from scratch (where the full table is represented by a single item). It consists of a set of itemEvents, one for each "key". Each one of these itemEvents is progressively modified or extended by "UPDATE" itemEvents that come from the supplier Data Adapter. The set itself is progressively extended or shortened by the "ADD" or "DELETE" itemEvents coming from the Data Adapter. If the itemEvents in the set are sent with the "ADD" command, they allow the Client to reconstruct the list.

Note that the snapshot is not a simple playback of the commands received so far, but an optimization of them. Moreover, the order of the keys involved in the snapshot itemEvents is unspecified. This should be taken into consideration, in particular, when unfiltered dispatching is desired.

When the Preprocessor implements the subscription to a new item toward the supplier Data Adapter, if the snapshot management for the item is supported by the Data Adapter, then the Data Adapter must use an asynchronous call-back to immediately send the snapshot, before any updates. However, if the snapshot for the item is not available (because the Data Adapter does not support it), the Preprocessor behaves as though the state of the item at the time of the subscription was empty.

The Client can indicate, at the moment of the subscription to a particular item (provided that it is of the type MERGE, DISTINCT or COMMAND), whether it intends to receive the snapshot first or to receive only the updates (as a default, the snapshot is not requested). In order to facilitate the handling of this differentiation, when the Preprocessor implements the subscription to a new item toward the supplier Data Adapter, it needs to distinguish the snapshot from successive updates utilizing a flag in the asynchronous call-back.

The snapshot returned by the Data Adapter should be calculated in the same way as the Preprocessor does. If this is not the correct way to determine the initial state of the item, then the snapshot for this particular item is not supported by Lightstreamer Server; the snapshot should not be requested by the clients and should not be supported by the Data Adapter at all. A typical situation of this kind is an item that represents an ordered list in which the new position of a "key" is carried by a field (it was described in the previous discussion of the unfiltered COMMAND mode); for such an item, the order and the contents of the received "ADD", "UPDATE" and "DELETE" events is important to maintain the ordering of the "keys" within the snapshot, while Lightstreamer Server Preprocessor ignores these information.

### 3.4.2 Prefilter

Each ItemEventBuffer implements an event filter, as a function of the assigned maximum frequency. In cases where the external Data Provider generates events with a frequency far greater than that sent to the Client, *only for the items handled in MERGE or DISTINCT mode with filtered dispatching* it is possible to implement a Prefiltering of the data at Preprocessor level, common to all users, in order to avoid a heavy task of filtering of each single ItemEventBuffer. In this way each ItemEventBuffer applies its own filtering policy on a set of data already "refined", with a notable increase in server performance.

The **MERGE Prefilter** works like an ItemEventBuffer with a size of 1 and of MERGE type, characterized by a maximum send frequency (**preFilterFrequency**).

The **DISTINCT Prefilter** works like an ItemEventBuffer with a size of 1 and of DISTINCT type, characterized by a maximum send frequency (**preFilterFrequency**). The dimension is kept at 1 in order to avoid delays on the sending of data.

The decision whether to apply or not the Prefiltering for each item is up to the Metadata Adapter, as well as the choice for preFilterFrequency.

NOTE: if a preFilterFrequency for an item is set, then *the item cannot be subscribed with unfiltered dispatching* in MERGE or DISTINCT mode; it can, however be subscribed in RAW mode.

### 3.4.3 Preprocessor and Modes

The following table summarizes the allowed behaviors of the Preprocessor (with respect to snapshot management, prefiltering, and filling) for the different subscription modes.

Mode	Allow Snapshot	Allow Prefilter	Allow Filling
RAW	No	No	No
MERGE	Yes	Yes	Yes
DISTINCT	Yes	Yes	No
COMMAND	Yes	No	Yes

Table 3 - Allowed behaviors of the Preprocessor.

**Filling** means that the Preprocessor will fill any empty field coming from the Data Adapter with the latest available value for such field.

Filling is a logical process. It does not imply that when any field is unchanged across subsequent updates in MERGE mode, each time the same field will be redundantly delivered to the Client. The network protocol of Lightstreamer is highly optimized and the whole system guarantees the data coherency while saving on used bandwidth.

### 3.5 Roles in the Choice of Parameters

Different actors contribute in the choice of the previously discussed parameters. For each session the actors involved in the choice are the Client, the Kernel, the Data Adapter and the Metadata Adapter. The following table summarizes the different roles of the actors in the choice of each parameter.

<i>Parameters</i>	<b>Client</b>	<b>Kernel</b>	<b>Data Adapter</b>	<b>Metadata Adapter</b>
<b>maxBandwidth</b>	can ask	-	-	must decide
<b>maxFrequency</b>	can ask (only for MERGE, DISTINCT, COMMAND modes)	Can limit furtherly, based on license terms	-	must decide
<b>unfiltered dispatching</b>	can ask (only for MERGE, DISTINCT, COMMAND modes)	must refuse the request, if not allowed	-	must decide
<b>Mode</b>	must ask	must refuse the request, if not allowed	-	must decide
<b>bufferSize</b>	can ask (only for MERGE, DISTINCT)	need not follow the decision	-	must decide
<b>Snapshot</b>	can ask (only for MERGE, DISTINCT, COMMAND modes)	-	should collaborate by sending the initial snapshot	-
<b>Snapshot length in DISTINCT mode</b>	can ask (only for DISTINCT mode)	-	should collaborate by sending enough data for the initial snapshot	must decide the DISTINCT history length
<b>preFilterFrequency</b>	-	-	-	must decide

**Table 4 - Choiche of parameters.**

**NOTE:** The Client can ask for different bandwidth and item frequency at a later time; the request is reexamined by the Metadata Adapter.

**NOTE:** Metadata Adapter decisions can be overridden through JMX settings. Then, in case of a request for different bandwidth or item frequency by the Client, the last JMX settings are used for the decision.



## 4 The Adapters

---

Each Lightstreamer session requires the presence of an Adapter Set, which is made up of one Metadata Adapter and one or multiple Data Adapters. The roles of these two types of Adapters are very different and enable for a good separation of concerns.

A Data Adapter only knows about individual items. It has no knowledge about Subscriptions, Item Groups, and Field Schemas. Furthermore, it has no notion at all of users.

A Metadata Adapter knows about Subscriptions, users, Item Groups and Field Schemas. It has a very fine-grained control over the life cycle of user sessions, including authentication, authorization and entitlement.

This chapter shows some sequence diagrams to help clarify the typical workflows of the interactions between Lightstreamer Server (the Kernel) and a Metadata Adapter, and between Lightstreamer Server (the Kernel) and a Data Adapter.

### 4.1 The Metadata Adapter

---

Figure 8 illustrates the typical workflow generated by a user's activities with respect to the session-related Metadata Adapter, from authentication (session start) to log off (session end). (Adapter initialization is not shown).

The complete specification of the methods invoked bidirectionally between the Server and the Metadata Adapter is available under the form of "API reference" within the Adapter SDKs. The interactions between the Client and the Server use a logical notation and does not correspond directly to the physical network protocol employed.

**Authentication.** When creating a new Lightstreamer session, the Client can pass the Server some credential information (two generic strings can be used, conventionally named "user" and "password"; in addition, cookies, custom HTTP headers, and client-side SSL/TLS certificates can be used). The Server validates the credentials against the Metadata Adapter, that throws a specific exception if the validation fails. If the user is authenticated, a LS\_session token is sent back to the Client, that will be used for all subsequent interactions in order to identify the session.

**QoS, Authorization, Entitlement.** The Metadata Adapter is responsible for assigning the quality of service to each user (in terms of bandwidth, frequency, etc.). It also interprets the "metadata" coming from the Client (Item Group and Field Schema names) and transform them into a form that can be used with the proper Data Adapter (that is, a list of item names and a list of field names). Furthermore, for each request coming from the Client, it should grant it or deny it. This applies to each subscription request and to the properties of each subscription (subscription mode, buffering, etc.).

**Log off.** The Metadata Adapter is notified when a session is closed (usually due to the user closing the client application).

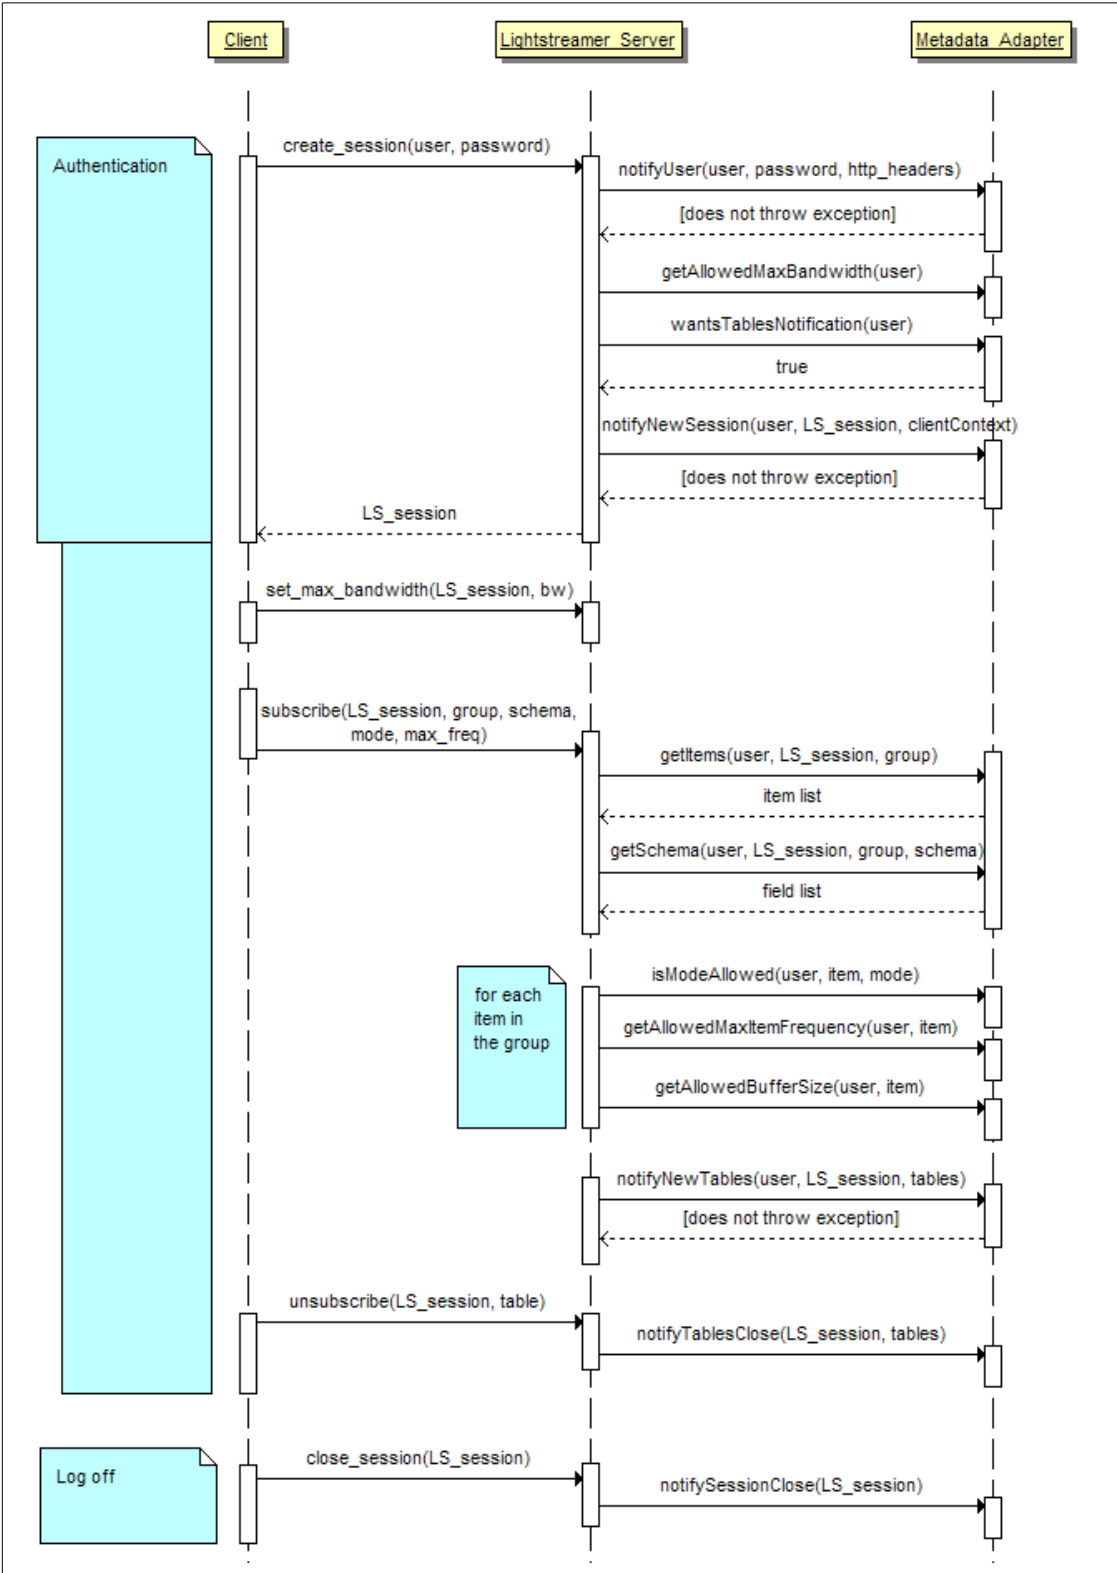


Figure 8- A typical Metadata Adapter workflow for a user.

Figure 9 focuses on the **authorization** process and illustrates the best practice used in typical scenarios. The actual authentication is usually handled by a legacy application server, irrespective of Lightstreamer. Some sort of `session_id` is sent back to the Client (through cookies, response payload or any other technique). When the Client creates the Lightstreamer session, instead of sending again the full credentials (usually involving a password) to Lightstreamer Server, it sends just the `session_id` (or the username and the `session_id`). The Metadata Adapter is passed this information and validates the `session_id` against the application server that generated it (or a database or whatever back-end system).

The advantage of this practice is that the user's password is never used outside the legacy system. Also, since the communication between the Client and Lightstreamer Server can leverage HTTPS, the `session_id` can be encrypted, so that it cannot be intercepted on the network.

This is just an example, because the Metadata Adapter is usually implemented in a custom way for each project, when integrating Lightstreamer into an existing or a new system.

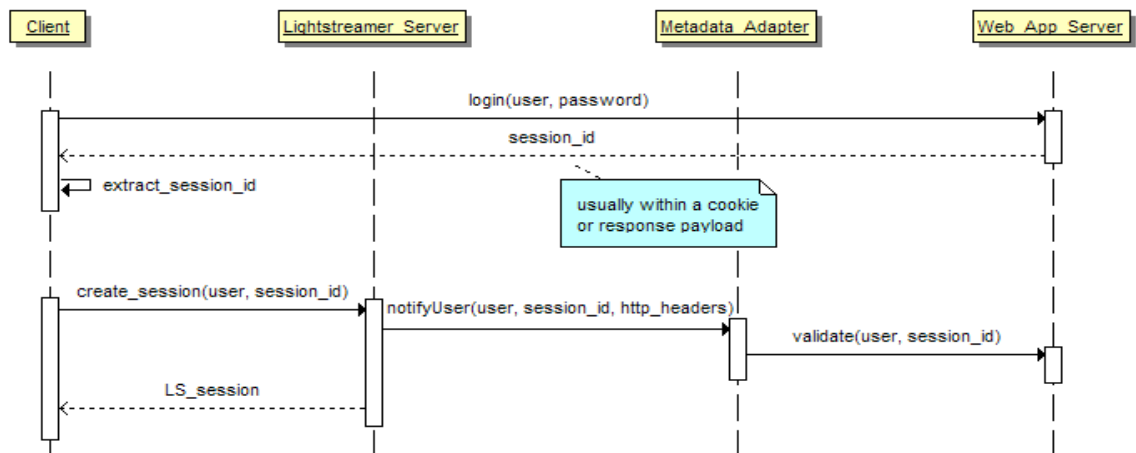


Figure 9- Authentication through the Metadata Adapter and an external server.

#### 4.1.1 Authentication and Authorization Examples

A simple example showing authentication and authorization in action is available online:

- Adapter code:  
<https://github.com/Lightstreamer/Lightstreamer-example-AuthMetadata-adapter-java>
- Client Code and live demo:  
<https://github.com/Lightstreamer/Lightstreamer-example-AuthMetadata-client-javascript>

## 4.2 The Data Adapter

Figure 10 illustrates the typical workflow with respect to a Data Adapter. At Server's start-up the Data Adapter is initialized and a listener is passed for all subsequent asynchronous callbacks (that is, from the Data Adapter to the Kernel). When a new subscription targeted to this Data Adapter is received by the Server, after extracting the individual items through the session-related Metadata Adapter (see the previous section), the Server activates the new item subscription on the Data Adapter. Before that, the Server needs to know if the Adapter is able to provide an initial snapshot for an item. After the subscription has been activated, the Data Adapter will start sending back asynchronous events. The Server will use these updates to generate events for the Client. When no more Clients are subscribed to a given item, that item is unsubscribed from the Data Adapter too.

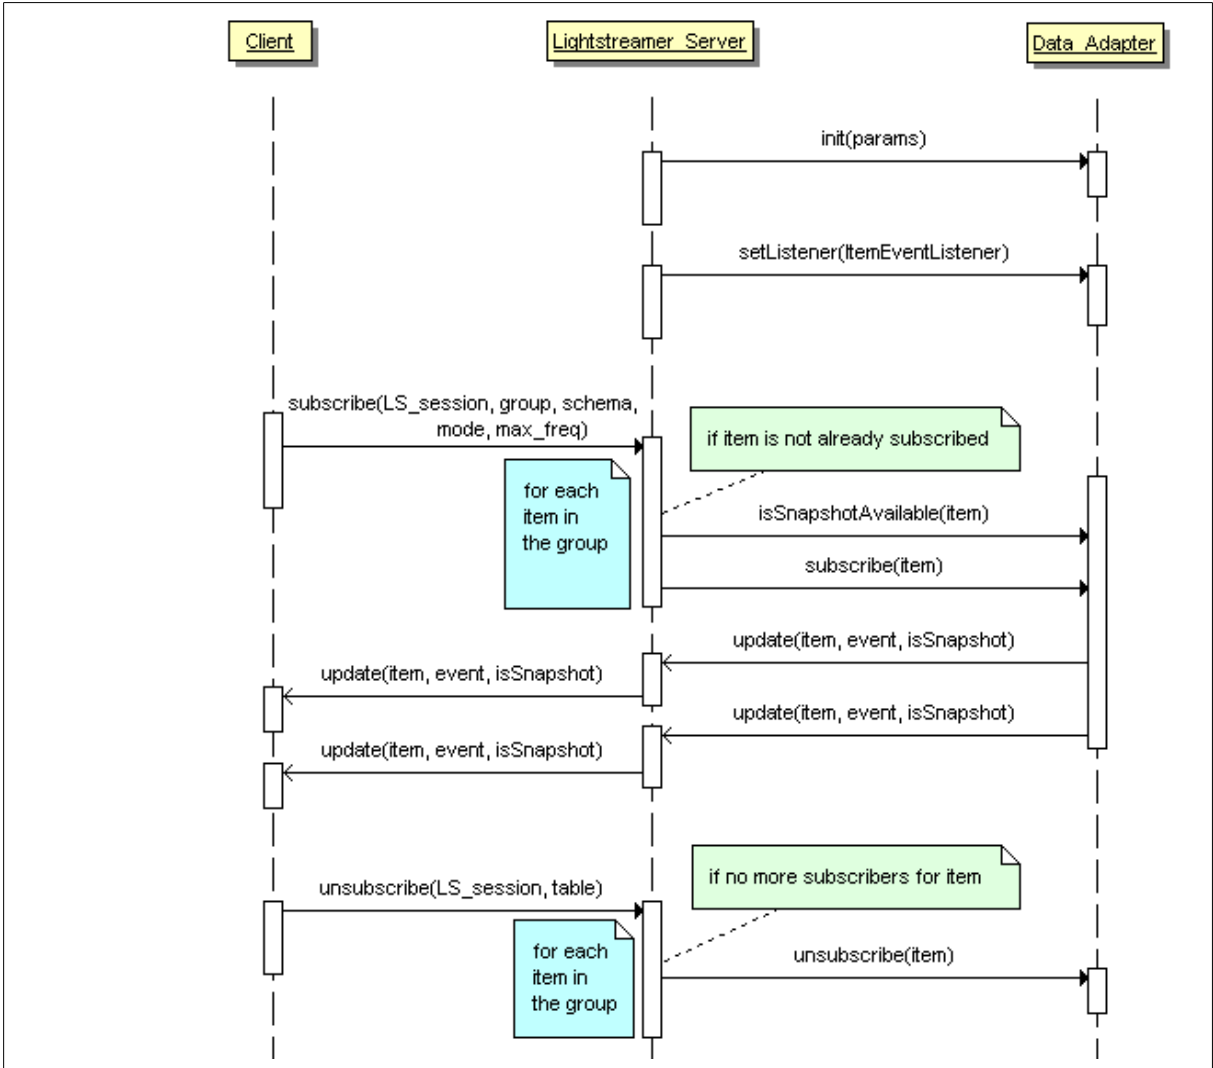


Figure 10- Data Adapter general workflow.

Figure 11 shows a concrete example. Two Clients are involved. Client B subscribes to item "item10", requesting a schema comprised of the fields "bid" and "ask". The Server will start serving that item after activating the subscription on the Data Adapter. Notice that the Data Adapter, as a first event, sends back a full snapshot for "item10". When Client A subscribes to "item10", the Server is able to serve it (including the initial snapshot) without asking anything to the Data Adapter. At this point the Server delivers the updates on "item10" to both the Clients. When both the Clients unsubscribe from "item10", the unsubscription is propagated to the Data Adapter, that should stop injecting events pertaining to "item10". When a Client subscribes to "item10", all the process starts again from scratch, that is, the Data Adapter receives the subscription and must send back the initial snapshot, followed by the real-time updates.

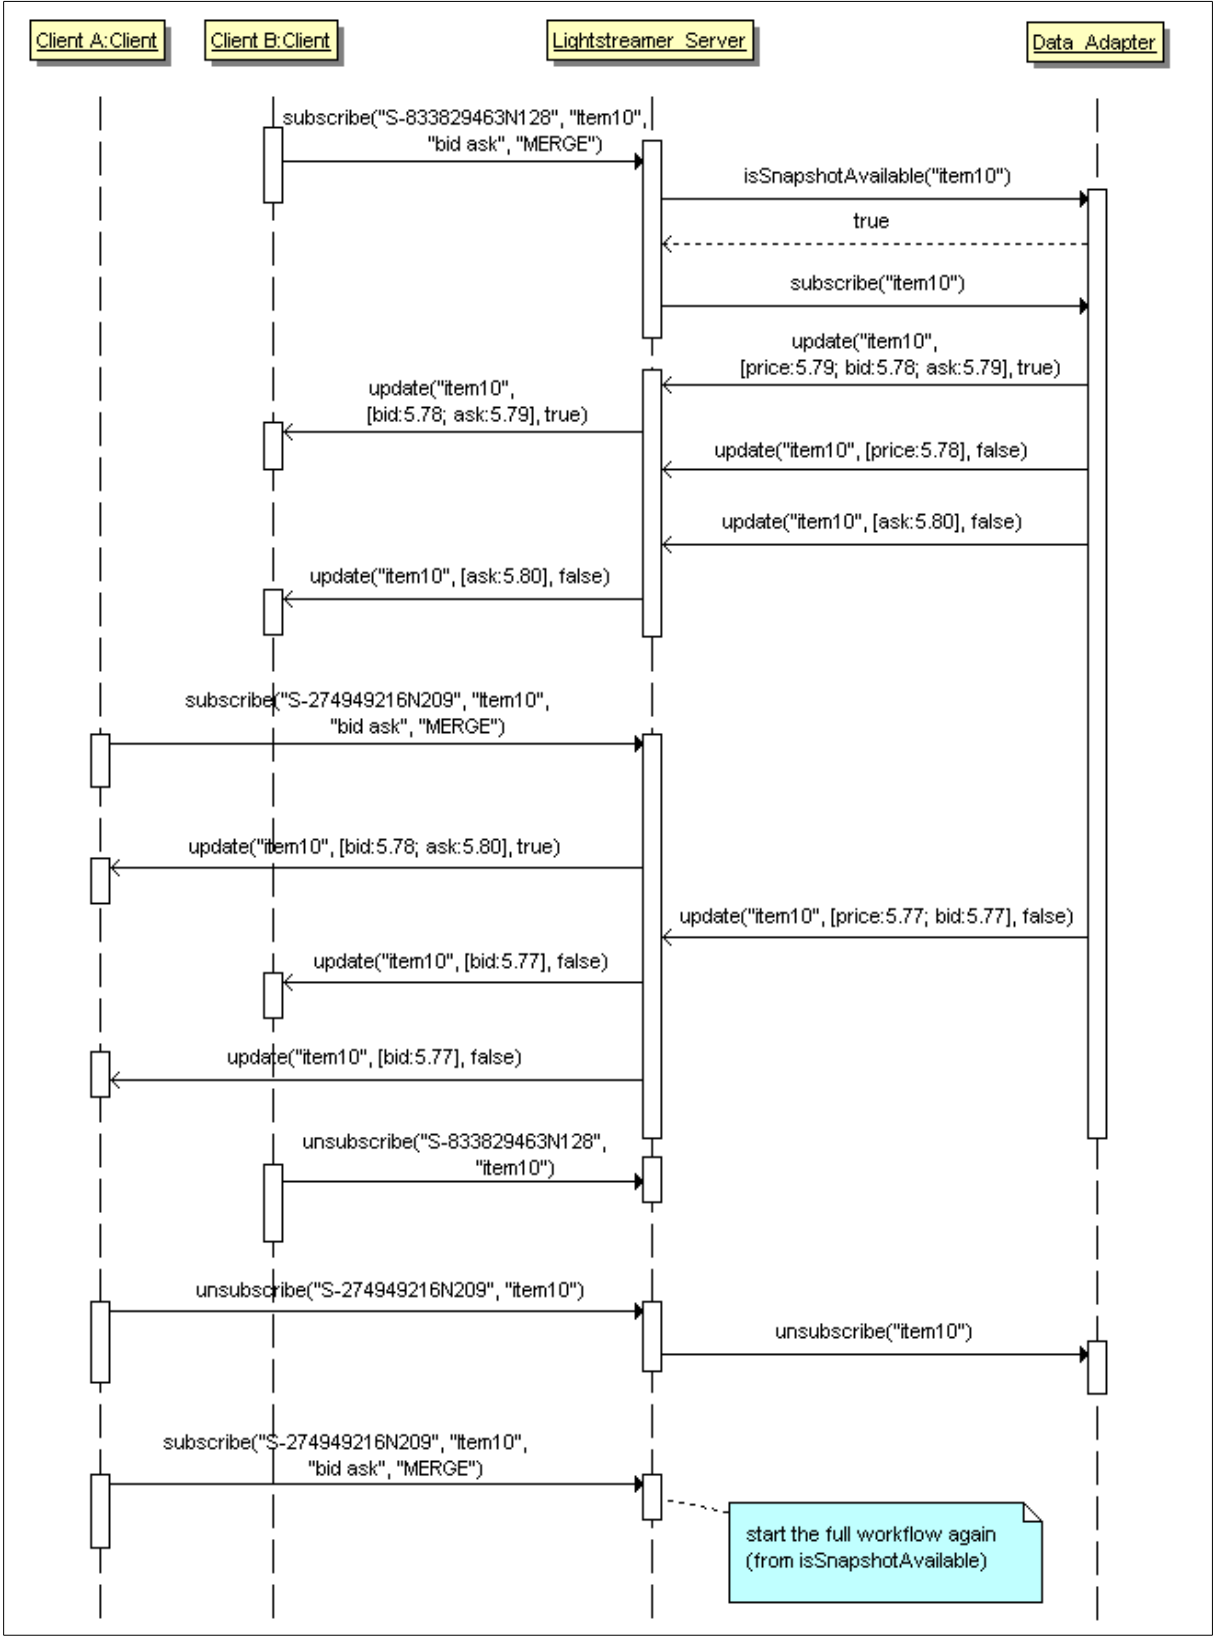


Figure 11- A concrete scenario of a Data Adapter workflow

#### 4.2.1 An Internal Data Adapter: The Monitoring Data Adapter

---

The Server provides a special Data Adapter, the monitoring Data Adapter, which can be included in one or multiple custom Adapter Sets, by specifying the reserved "MONITOR" name in place of the Data Adapter class in the Adapter Set configuration.

This allows the clients to request internal Server statistics and monitoring data along with other application data. This is the same Data Adapter used by Monitoring Dashboard to provide real-time statistics on server's internals.

The custom Metadata Adapter included in the Adapter Set should identify requests directed to the monitoring Data Adapter and operate any desired access restrictions to these data.

The following items are supplied:

- **monitor\_statistics**

Internal load statistics, provided in MERGE mode without snapshot support.

The available fields are the following:

- **MEMORY.TOTAL** The total memory (expressed in Bytes) allocated for the heap.
- **MEMORY.FREE** The free memory (expressed in Bytes) in the heap.
- **THREADS** The total number of threads in the Java virtual machine.
- **POOL.SIZE** The global number of threads belonging to one of the thread pools used by Lightstreamer Server to manage the various processing stages. This value is the sum of POOL.ACTIVE and POOL.WAITING, but for possible further threads currently in a transition phase.
- **POOL.ACTIVE** The total number of threads belonging to one of the thread pools used by Lightstreamer Server which are currently executing a task.
- **POOL.WAITING** The total number of threads belonging to one of the thread pools used by Lightstreamer Server which are currently unused and waiting for a new task to be enqueued.
- **POOL.QUEUE** The total number of tasks waiting to be processed by any of the thread pools used by Lightstreamer Server.
- **POOL.QUEUE\_WAIT** The average waiting time in milliseconds experienced by the tasks scheduled for processing by any of the thread pools used by Lightstreamer Server.
- **NIO.WRITE\_QUEUE** The number of socket write operations currently waiting to be selected by the NIO selector.
- **NIO.WRITE\_QUEUE\_WAIT** The waiting time, in milliseconds, averaged over all the socket write operations currently waiting to be selected by the NIO selector.
- **NIO.WRITE\_SELECTORS** The number of selectors and related threads currently used to handle write operation tasks.
- **NIO.TOTAL\_SELECTORS** The total number of selectors and related threads currently in use.
- **CLIENTS.CONNECTIONS** The number of current Connection sockets; it comprises Stream, Polling, Control, Web Server and unrecognized requests.
- **CLIENTS.MAX\_CONNECTIONS** The maximum number of concurrent Connection sockets reached in the life of the Server; it comprises Stream, Polling, Control, Web Server and unrecognized requests.

- **CLIENTS.SESSIONS** The number of current Sessions. Prestarted Sessions are not considered for this and the following session statistics. By **prestarted** Session we mean a session for which no related requests have been received yet. Binding a session with a new connection also gets it out of the prestarted state.
- **CLIENTS.MAX\_SESSIONS** The maximum number of concurrent Sessions reached in the life of the Server.
- **CLIENTS.STREAMING\_SESSIONS** The number of current Sessions that are associated with a Stream Connection.
- **CLIENTS.POLLING\_SESSIONS** The number of current Sessions that are associated with Polling Connections.
- **CLIENTS.ITEM\_SUBSCR** The total number of currently active Item subscriptions to Items. The value is cumulated over all client sessions, hence a single item, subscribed by multiple clients, ore even multiple times by the same client, will be counted multiple times.
- **CLIENTS.MAX\_ITEM\_SUBSCR** The maximum number of concurrently active subscriptions to Items reached in the life of the Server.
- **ITEMS.TOTAL** The total number of currently subscribed Items, covering all the Data Adapters from all the Adapter Sets plugged in the Server.
- **ITEMS.EVENTS\_SEC** The frequency of inbound events flowing from all Data Adapters, within any Adapter Set, into LS Kernel in the last sampling period. It is expressed as events per second.
- **ITEMS.FILTERED\_EVENTS\_SEC** The frequency of outbound events flowing from the Preprocessor of LS Kernel in the last sampling period. It is expressed as events per second.
- **UPDATES.EVENTS\_SEC** The frequency of outbound events flowing towards the clients. The value is cumulated over all client sessions, hence a single event produced by the Data Adapter could trigger multiple updates on one or multiple sessions. It is expressed as updates per second.
- **UPDATES.TOTAL\_IN** The total (cumulative) number of inbound events flowing from all Data Adapters, within any Adapter Set, into LS Kernel.
- **UPDATES.TOTAL\_OUT** The total (cumulative) number of outbound events flowing towards the clients. The value is cumulated over all client sessions, hence a single event produced by the Data Adapter could trigger multiple updates on one or multiple sessions.
- **UPDATES.TOTAL\_LOST** The total (cumulative) number of updates lost. Events can only be lost because of buffer size restrictions posed for safety purpose. This applies to Items subscribed to in RAW mode, or in any mode with unfiltered dispatching specified. This also applies to items subscribed to in COMMAND mode with filtered dispatching, restricted to "ADD" and "DELETE" events only (note that those events can also be filtered through matching). Any lost event is also notified to the client.
- **BANDWIDTH.TOTAL\_BYTES** The total (cumulative) number of bytes sent in the life of the Server (for Stream and Polling Connections only). The count includes all HTTP contents sent.
- **BANDWIDTH.CURRENT** The current global outbound bandwidth used by the Server (for Stream and Polling Connections only) in the last sampling period. It is expressed as Kilobits per second and it is comprised of network overheads.
- **BANDWIDTH.HIGHEST** The maximum global outbound bandwidth used by the Server in its life (for Stream and Polling Connections only). It is expressed as Kilobits per second and it is comprised of network overheads.



- **BANDWIDTH.ALLOCATED** The current outbound allocated bandwidth, calculated as the sum of the maximum outbound bandwidths allocated by the Server for each Session. It is expressed as Kilobits per second and it is comprised of network overheads.
- **BANDWIDTH.CLIENT\_AVERAGE** The average outbound bandwidth used by the Server for each Session (in the last sampling period). It is expressed as Kilobits per second and is comprised of network overheads.
- **MESSAGES.EVENTS\_SEC** The frequency of Client Messages submitted to the Metadata Adapter in the last sampling period. It is expressed as messages per second.
- **MESSAGES.TOTAL\_HANDLED** The total (cumulative) number of Client Messages submitted to the Metadata Adapter in the life of the Server.
- **MESSAGES.TOTAL\_BYTES** The total (cumulative) number of bytes submitted to the Metadata Adapter in the life of the Server and related to Client Messages associated to sendMessage requests. The count includes the byte length of the String objects carrying the messages, according with their internal UTF-16 encoding.
- **MESSAGES.CURRENT\_THROUGHPUT** The current global throughput (measured in the last sampling period) related to the submission to the Metadata Adapter of Client Messages associated to sendMessage requests. It is expressed as Kilobits per second and it counts the byte length of the String objects carrying the messages, according with their internal UTF-16 encoding.
- **MESSAGES.HIGHEST\_THROUGHPUT** The maximum global throughput (occurred in the life of the Server) related to the submission to the Metadata Adapter of Client Messages associated to sendMessage requests. It is expressed as Kilobits per second and it counts the byte length of the String objects carrying the messages, according with their internal UTF-16 encoding.
- **PUSH\_NOTIFICATIONS.DEVICES** The number of Devices currently served by the MPN Module, when enabled. Actually, multiple apps can request Push Notifications for the same physical device; here, by Device we mean the combination of a physical device and an app.
- **PUSH\_NOTIFICATIONS.MAX\_DEVICES** The maximum number of Devices concurrently served by the MPN Module reached in the life of the Server.
- **PUSH\_NOTIFICATIONS.NOTIFICATIONS\_SEC** The frequency of Push Notifications submitted by the MPN Module to the proper external service for delivery in the last sampling period. It is expressed as notifications per second.
- **PUSH\_NOTIFICATIONS.TOTAL\_SENT** The total (cumulative) number of Push Notifications submitted by the MPN Module to the proper external service for delivery in the life of the Server.
- **PUSH\_NOTIFICATIONS.TOTAL\_BYTES** The total (cumulative) number of bytes related with Push Notifications submitted by the MPN Module to the proper external service for delivery in the life of the Server. The count only includes the notification payload.
- **PUSH\_NOTIFICATIONS.CURRENT\_THROUGHPUT** The current global throughput (measured in the last sampling period) related with Push Notifications submitted by the MPN Module to the proper external service for delivery. It is expressed as Kilobits per second and it only counts the notification payload.
- **PUMP\_WAIT.SLEEP** The average extra-time experienced for sleep() system calls. It is expected that the difference between the expected time and the observed time increases when the available CPU decreases, though significant delays with plenty of CPU available have sometimes been observed.

- **PUMP\_WAIT.NOTIFY** The average time experienced between a notify() system call and the exit from the related wait() in a producer/consumer interaction. It is expected that this time difference increases when the available CPU decreases, though significant time differences with plenty of CPU available have sometimes been observed.

- **monitor\_identification**

Basic data about the Server, provided in MERGE mode (but they only consist of a snapshot).

The available fields are the following:

- **EDITION** The edition in which Lightstreamer Server is running.
- **VERSION** The version of Lightstreamer Server.
- **LICENSE\_TYPE** The license type used to run Lightstreamer Server: "FREE", "DEMO", "LICENSE\_FILE" or "LICENSE\_SERVER".
- **CLIENT\_ID** The client id of the license used to run Lightstreamer Server. For a "DEMO" or "FREE" license type, "-" is returned.
- **LOCAL\_IP** The local IP address of the host running Lightstreamer Server.
- **LOCAL\_HOST** The local hostname of the host running Lightstreamer Server.
- **STARTUP\_TIME** The timestamp representing the number of milliseconds elapsed between 1 January 1970 00:00:00 UTC and the moment the Lightstreamer Server instance was started.

- **monitor\_socket\_<name>**

Information about a server socket, as configured through the <http\_server> or <https\_server> elements, recognized through the specified name. The data is provided in MERGE mode (but it only consists of a snapshot).

The available fields are the following:

- **SERVER\_NAME** the name of the server socket, as configured through the "name" attribute.
- **LOCAL\_PORT** the configured value for the local interface IP address on this server socket, or null if no configuration has been supplied.
- **LOCAL\_INTERFACE** the local port where this server socket is listening.
- **IS\_SSL** "true" if the server socket is for https connections.

Note that configuring the same name for different server sockets is allowed; in this case, however, only data for one of the configured server sockets will be returned for the related item.

- **monitor\_log\_error, monitor\_log\_warning, monitor\_log\_info, monitor\_log\_output**

The log produced by the Server on the preconfigured "LSProducer" appender, provided in DISTINCT mode without snapshot support. The different items select messages based on the priority level, in the following way:

- monitor\_log\_error: ERROR and FATAL
- monitor\_log\_warning: WARN
- monitor\_log\_info: INFO
- monitor\_log\_output: all messages

The available fields are the same for all these items; they are the following:

- **TIME** a timestamp in the "HH:mm:ss" form.

- **TIME\_MS** a UTC timestamp.
- **CLIENT.IP** the IP address of the related client (only available for some client-related INFO level messages).
- **CLIENT.NAME** the hostname of the related client, if the hostname lookup is enabled and successful; the IP address otherwise (only available for some client-related INFO level messages).
- **MESSAGE** the log message.
- **THREAD** the name of the thread originating the log message.
- **COUNTER** a progressive counter, started on item subscription.

## 5 Mobile Push Notifications

Lightstreamer supports native push notifications on **Google Cloud Messaging** (GCM) for Android devices and **Apple Push Notification Service** (APNs) for iOS devices.

Mobile push notifications (MPN for short) are handled by a special Lightstreamer module called **MPN Module**. Since MPNs are designed to reach the user even when their application is not running, they can't be bound to a common Lightstreamer session. For this reason, the MPN Module has the ability to manage special subscriptions, called MPN subscriptions, which survive the scope of a session and are persisted on an SQL database.

Thanks to the MPN Module, items that are already managed by the Data Adapters can now be pushed to mobile clients through native push notifications, in addition to Web based protocols. Making items available via push notifications does not require any further server-side development.

These MPN subscriptions are kept alive by the MPN Module and restarted each time the Lightstreamer Server restarts. Multiple instances of the Server (e.g. in a clustered configuration) are able to pass MPN subscriptions to each other by leveraging the database. When an MPN subscription produces an update, this is formatted in an MPN notification according to user's specifications and sent to the appropriate MPN provider (APNs or GCM). An extended view of the system architecture, considering the MPN Module, is the following:

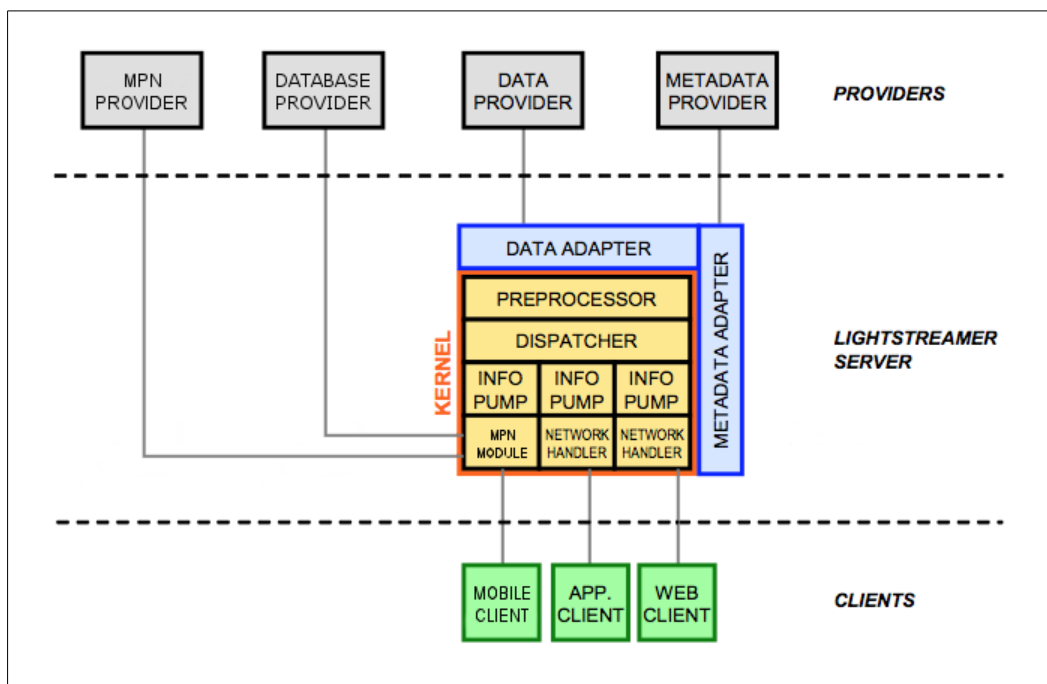


Figure 12-Extended view of the architecture with MPN Module

Mobile clients can activate MPN subscriptions by using special APIs provided by their Lightstreamer SDK. An MPN subscription activation requires an active Lightstreamer session and takes both the parameters of a common subscription (data adapter, group ID, schema, etc.) and the additional parameters needed to specify how (and when) the notifications must be formatted and delivered.

The MPN Module is configured through the `<mpn>` section of the Lightstreamer Server configuration file.

**Live examples with full source code** are available online, to get started with Mobile Push Notifications very easily:

→ **iOS client example:**

<https://github.com/Lightstreamer/Lightstreamer-example-MPNStockList-client-ios>

→ **Android client example:**

<https://github.com/Lightstreamer/Lightstreamer-example-MPNStockList-client-android>

## 5.1 MPN Subscriptions

MPN subscriptions are activated by clients using specific APIs made available on certain SDKs. Activating an MPN subscription requires an active session, even if the subscription itself is not bound to the session, as introduced earlier. Nevertheless, a session is necessary so that the Metadata Provider may validate and authorize the MPN subscription request. To this purpose, an MPN subscription request can only refer to items served by the Adapter Set associated with the hosting session. By the way, note that in case of a clustered configuration, a subscription may be passed from one instance to another; this leans on the fact that the mounted Adapter Sets must be identical on all instances.

MPN subscriptions are subject to limitations, due to their nature of delivering updates directly to the end user (in most cases). As introduced before, an MPN subscription is made of both the parameters of a common subscription and a few more to specify how and when the notification must be delivered. On the common subscription parameters, the following limitations apply:

- only **MERGE** and **DISTINCT** modes are allowed;
- unfiltered dispatch can't be used;
- snapshot can't be used;
- selectors and customizers can't be used (but the additional parameters dedicated to MPN may better serve the same purposes, see below).

Additional parameters depend on the specific MPN platform, but some general rules apply:

- one or more parameters are dedicated to *how to format* the notification message: these parameters accept a special formatting syntax that lets you compose the message with the content of the update (see §5.1.1);
- all MPN platforms accept a new parameter called **trigger expression** that lets you specify *when to send* the notification message (see §5.1.2).

MPN subscriptions are also subject to frequency limitations: most MPN Providers don't allow a high pace of notifications to be sent to the device. If notifications are sent with a too high pace the provider may forcibly disconnect the MPN Module and ban it as a client.

For this reason a specific maximum frequency is set in terms of minimum delay between subsequent notifications. This limit is per MPN Provider and may be configured in the specific MPN Provider's configuration file (see §5.3). The limit is enforced on a per device-app pair basis, i.e.: two different apps on the same device are counted separately towards the frequency limit.

### 5.1.1 The Formatting Syntax

---

All parameters dedicated to formatting the notification message accept the embedding of field values, coming from the update, by using a named-argument or indexed-argument specification:

- a named argument is specified by its field name surrounded by `#{` and `}`, e.g.: `"#{last_price}"`;
- an indexed argument is specified by its 1-based field position surrounded by `#[` and `]`, e.g.: `"#[1]"`.

Thus, a format specified with "The value of stock `#{stock_name}` is now `#{last_price}`" will be composed into something like "The value of stock Anduct is now 13.63" when delivered to the end user.

Consider that the schema and group ID of the subscription are validated through the Metadata Provider as usual. For this reason, the use of named arguments is possible only if the Metadata Provider is a subclass of `LiteralBasedProvider` or an equivalent class, while the use of indexed arguments is always possible.

### 5.1.2 The Trigger Expression

---

In many cases the delivery of a notification for each update is not desirable. Instead, the deliver of a notification when a specific condition is met may be preferable, e.g.: *send me an update when the stock value rises above 500.0*. Trigger expressions are designed for this purpose.

A trigger expression is a **Java boolean expression** that is evaluated against the update values and, if found true, triggers the send of the notification. An MPN subscription actually works differently whether a trigger expression is present or absent:

- in presence of a trigger expression, the MPN subscription delivers the notification only when the expression evaluates to true the first time, no other notifications are sent after that;
- In absence of a trigger expression, the MPN subscription delivers a notification each time the subscription produces an update.

The trigger expression may contain named arguments and indexed arguments (see §5.1.1), and from a Java point of view they must be considered as *String* variables. So, they must undergo appropriate type conversion if a comparison is needed, e.g.:

- `"#{stock_price} > 500.0"` is an **invalid** trigger expression, since you can't compare a `String` with a number;
- `"Double.parseDouble(#{stock_price}) > 500.0"` is a **valid** trigger expression, as it correctly converts the `String` to a number before comparison.

Named and indexed arguments are converted to variables before the trigger expression is evaluated. Specifically, these variables have the prefix `LS_MPN_field` and a suffix which corresponds to the field index, e.g.: `"Double.parseDouble(#{stock_price}) > 500.0"` is converted in `"Double.parseDouble(LS_MPN_field3) > 500.0"` if `stock_price` is the 3rd field in the schema.

For security reasons, trigger expressions may make use of any Java system class (by using their fully qualified name) but no other classes. You can't access the Lightstreamer Server JMX APIs, for example, or your own provided  *jars*. However, distinct executions of trigger expressions may share the same static context and perform side-effects on this context; this is not limited to the same underlying device.

For this reason, the Metadata Adapter has the opportunity to check any submitted trigger and it is responsible for the acceptance of the trigger expression. This is to avoid that a hacker may compromise the Server integrity with a maliciously crafted trigger expression. In simple cases, a prefiltering of the trigger expressions performed by the Kernel, and based on matching with regular expressions, can be leveraged, by specifying the regular expressions in the MPN Provider configuration file (see §5.3).

In particular, MPN provider configuration files contain a specific section **<trigger\_expressions>** for each application, where a list of (Java) regular expressions can be specified to validate triggers submitted by the clients. For a trigger to be accepted, it must match at least one regular expression.

For example, a list like the following:

```
<trigger_expressions>
  <accept>Double.parseDouble(\${[A-Za-z0-9_+]\} [!&#x26;] [+]?(\d*\.)?d+</accept>
  <accept>Double.parseDouble(\${[d+]\} [!&#x26;] [+]?(\d*\.)?d+</accept>
</trigger_expressions>
```

will let pass a trigger like "Double.parseDouble(\${stock\_price}) > 500.0", but of course not one like "System.exit(0)". This simple example shows why accept-all regular expressions like ".\*", even if possible, are strongly discouraged.

### 5.1.3 MPN Subscription ID and Coalescence

Once an MPN subscription has been activated, it is identified by a globally unique ID called *MPN key*. With this key, the MPN subscription may be inquired, modified or deactivated from any session (the one where it has been created or any successive one).

This means that if the typical workflow of Lightstreamer apps:

1. connect to Lightstreamer
2. activate the subscription
3. receive updates

is applied in the same way to MPN subscriptions:

1. connect to Lightstreamer
2. activate the MPN subscription
3. receive notifications

it will lead to multiple notifications for the same update. In fact, while real-time subscriptions are lost when the session is closed (i.e. when the app disconnects or is terminated), MPN subscriptions are not, they remain live on the Server. Each activation would then add *a new* MPN subscription.

The workflow for MPN subscriptions should then be changed in the following way:

1. connect to Lightstreamer
2. look up the MPN subscription's key in the app's storage:
  - if it is present, the MPN subscription is already active (and may be optionally inquired to obtain its details)
  - if it is not present, activate the MPN subscription and store its key
3. receive notifications

Alternatively, the MPN subscription may be activated using the **coalescing** flag. This flag has been introduced to simplify the workflow in cases where there is a fixed list of a few (or just one) MPN subscriptions. It may be specified on the MPN activation call, and its effect is that if the MPN subscription shares the same **base parameters** (i.e. adapter set, data adapter, group ID, schema and trigger expression) with an existing MPN subscription, the activation is actually considered as a modification of the existing MPN subscription. That is: just one MPN subscription is created per unique combination of these parameters, and subsequent activations just update the other non-base parameters (e.g. the notification format).

Using coalescing, the simplified workflow becomes:

1. connect to Lightstreamer
2. activate the MPN subscription with **coalescing = true**
3. receive notifications

Supposing base parameters are always the same, each activation will *not* lead to a new MPN subscription each time. There will be just one MPN subscription, and the user will not receive multiple notifications for the same update.

Of course, coalescing makes sense if always used on the same set of MPN subscriptions. A mixed use of this flag on subscriptions with the same base parameters may lead to the case of a coalescing operation with multiple active subscriptions eligible; in that case, Lightstreamer would pick one at random.

## 5.2 The Database Provider

---

Persistence in the MPN Module is managed through *Hibernate*, a well-known object-relational mapping library that efficiently encapsulates the underlying SQL database dialect and details. Hibernate has its own configuration file, which must be pointed to by the **<hibernate\_config>** tag in the **<mpn>** section of the Lightstreamer Server configuration file.

The Hibernate configuration file contains options for:

- the JDBC connection parameters;
- the specific SQL dialect to be used;
- the JDBC connection pool parameters;
- debugging options like whether SQL queries have to be dumped or not, and others.

The provided file is preconfigured for a MySQL database engine, while sample parameters are also shown for HSQL (suitable for testing) and Oracle databases.

The Hibernate configuration file also contains pointers to the database entities mapping files. These files, in turn, contain options to set column names and sizes, indexes, and so on. The provided default values are compliant with predominant databases, but they may have to be tuned in case of deployment specificities.

The database tables store the current state of the active notifications and the provided column sizes affect the length limits that will be imposed upon single fields of client activation requests. Modifying the column size configuration in order to change the limits is allowed, if supported by careful testing. The column meaning can be inferred from mapping files; on the other hand, a detailed description of the persisted data is currently not provided.



Once the Hibernate configuration file has been set, the appropriate libraries (e.g. the database's JDBC driver and its dependencies) must be added under the **lib/mpn/hibernate** directory of the Lightstreamer Server installation. The JDBC driver is typically distributed together with the database.

### 5.2.1 The Database in a Clustered Configuration

In case of a clustered configuration, *all Lightstreamer Server instances must point to the same database provider*. It is up to the Integrator to provide an SQL database with appropriate redundancy and load balancing.

Moreover, all the nodes of the cluster must share the same adapter sets. In fact, nodes may pass MPN subscription to each other during events like fail-overs or restarts: in these cases the MPN module must be able to start MPN subscriptions previously activated by another node.

In case of configuration changes with MPN subscriptions already in place (such as the adapter set changing name), the use of a custom SQL script to update existing rows may be necessary. The database structure has been designed to be inspectable and, when client APIs can't be used, updatable. An Integrator that finds themselves in this situation may contact [support@lightstreamer.com](mailto:support@lightstreamer.com) for advices on how to proceed.

### 5.2.2 DBMS Compatibility

The MPN Module has been successfully tested with the following DBMS:

DBMS	Compatible	Version tested	JDBC driver tested	Notes
<b>HSQldb</b>	✓	2.3.0	embedded	None.
<b>MySQL</b>	✓	5.5.35	5.1.28	None.
<b>Oracle</b>	✓	11g	ojdbc6 release 2	None.
<b>MS SQL Server</b>	✓	2008 R2	sqljdbc4 2008 R2	None.
<b>PostgreSQL</b>	✓	9.1.13	9.3-1101 jdbc4	None.
<b>IBM DB2</b>	✓	10.5 FP3	10.5 FP3	Requires a change in the table name for Module objects, may be set in the <b>Module.hbm.xml</b> file. Default name is "MODULES", which is a reserved keyword.

Other DMBS may result compatible or require some adjustments. Hibernate usually provides good compatibility with well-known DBMS.

## 5.3 The MPN Provider

At the time of writing, supported MPN providers are APNs (Apple Push Notification service), used by both the iOS and the macOS Lightstreamer SDKs, and GCM (Google Cloud Messaging), used by the Android Lightstreamer SDK. More providers will come with future releases.

### 5.3.1 APNs

---

The Apple Push Notification Service Provider, or *APNs notifier* for short, is configured through a specific configuration file pointed to by the **<apns\_notifier\_conf>** tag in the **<mpn>** section of the Lightstreamer Server configuration file.

This file is subdivided in **<app>** sections, with each section specifying parameters for a served mobile app. To be able to send notifications through APNs, each app must specify its own certificate (obtainable through Apple's Developer Portal) and related password, in **p12** format. Recall that Apple enforces the service level: development (sandbox) certificates may be used when debugging apps, but to use a production certificate you need a production (e.g. distribution or ad-hoc) provisioning profile for your app.

As said in §5.1.2, each app configuration may also specify a list of (Java) regular expressions to be used to validate trigger expressions. If omitted, no trigger will be accepted for the app.

### 5.3.2 GCM

---

The Google Cloud Messaging Provider, or *GCM notifier* for short, is configured through a specific configuration file pointed to by the **<gcm\_notifier\_conf>** tag in the **<mpn>** section of the Lightstreamer Server configuration file.

This file is subdivided in **<app>** sections, with each section specifying parameters for a served mobile app. To be able to send notifications through GCM, each app must specify its own API key (obtainable through Google's Developer Console). For testing purposes, each app may be configured for *dry run*, i.e. send notifications to Google but avoid their delivery to the device.

As said in §5.1.2, each app configuration may also specify a list of (Java) regular expressions to be used to validate trigger expressions. If omitted, no trigger will be accepted for the app.