

Striving for Versatility in Publish/Subscribe Infrastructures

Roberto S. Silva Filho and David F. Redmiles
{rsilvafi, redmiles}@ics.uci.edu

Department of Informatics
Donald Bren School of Information
and Computer Sciences
University of California, Irvine


Presented at SEM 2005 – September 4th, 2005

Outline

- Motivation
- Versatility definition
- Approach
- Architecture overview
- Case studies
- Conclusions and future work

Motivation

- The publish/subscribe communication style provides:
 - data, flow and timing decoupling between producers and consumers of information
 - content-based filtering and communication
- This mechanism is usually implemented by a logically centralized infrastructure
 - intermediates the communication between publishers and subscribers in a distributed setting.



Motivation (continuation)

- For such properties publish/subscribe middleware has been used in different application domains such as:
 - software monitoring, groupware, workflow management systems, software development and deployment, mobile applications and so on.
- This wide range of applications have required different sets of services from the publish/subscribe infrastructure such as:
 - Advanced event processing, guaranteed event delivery, transactions, event persistency, secure communication channels, authentication, mobility support and many others

Motivation (continuation)

- In order to implement a distributed event-driven application, two main alternatives exist:
 - Use existing publish/subscribe infrastructures:
 - Standardized one-size-fits-all solutions: CORBA-NS or JMS
 - Minimal content-based routers such as ELVIN, SIENA, HERALD
 - Build new specialized pub/sub system
 - example: CASSIUS, GEM, YEAST and others.

Motivation (continuation)

- Those strategies, however, suffer from a fundamental problem:
 - They are not **flexible** enough [c.f. Parnas] :
 - They are usually not designed for change and evolution,
 - Nor to be expanded and contracted to address specific application needs

- Which results in:
 - The need for direct source code modification of existing solutions (when available)
 - The implementation of additional features at the application level
 - the build of new pub/sub infrastructures
 - resulting in the proliferation of incompatible proprietary infrastructures that are costly to evolve and maintain

Versatility

- In other words, current publish/subscribe infrastructures are not versatile enough to support their use in different application domains.

- Our concept of versatility comprises a set of properties:
 - Support for Evolution
 - Extensibility – add new functionality to the existing set
 - Programmability – redefine software behavior
 - Reuse
 - Support for Variability (footprint configuration)
 - Static (build or design time)
 - Dynamic (runtime)
 - Usability
 - Considerations about workplace environment
 - Nielsen's attributes: learnability, efficiency, memorability, few errors and satisfaction.
 - Preserving middleware requirements of:
 - Scalability, interoperability, heterogeneity and communication

Our approach:

YANCEES, a versatile
publish/subscribe infrastructure

Approach main characteristics

Based on the use of extensible languages, plug-ins and filters

- combining language and infrastructure evolution
- with static and dynamic plug-in configurations

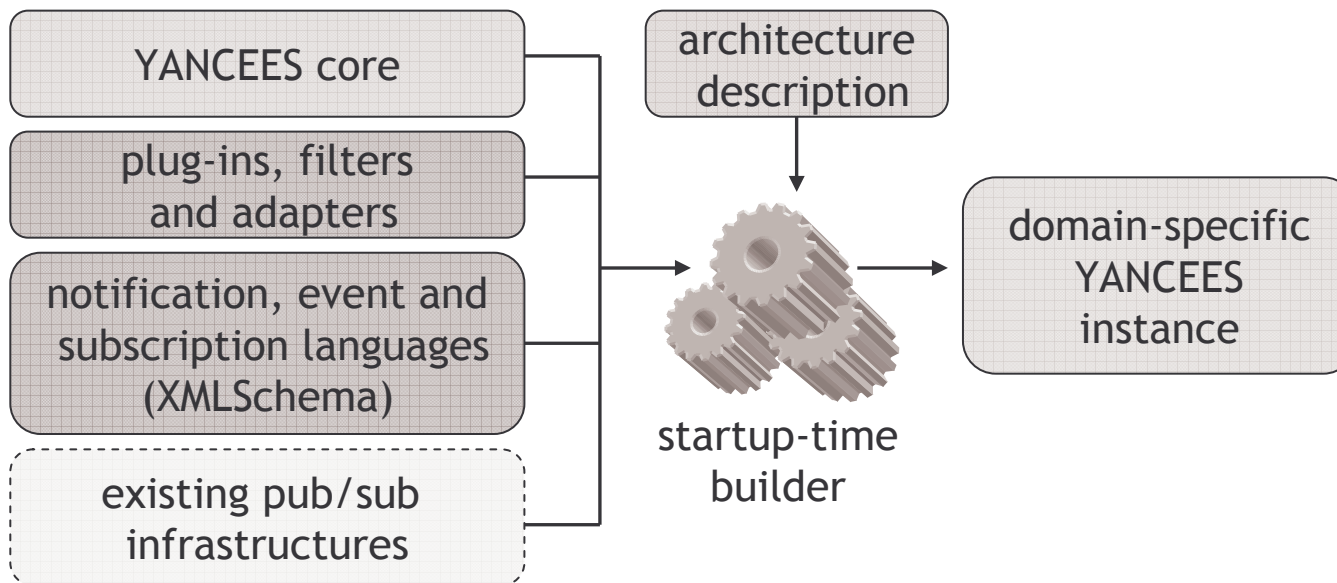
■ Built upon a micro kernel architecture style

- achieving interoperability and support for different event models and routing strategies

■ The architecture variability follows an extended version of Rosenblum and Wolf's [24] publish/subscribe design dimensions

■ The components are put together with the help of runtime parsers and static configuration managers

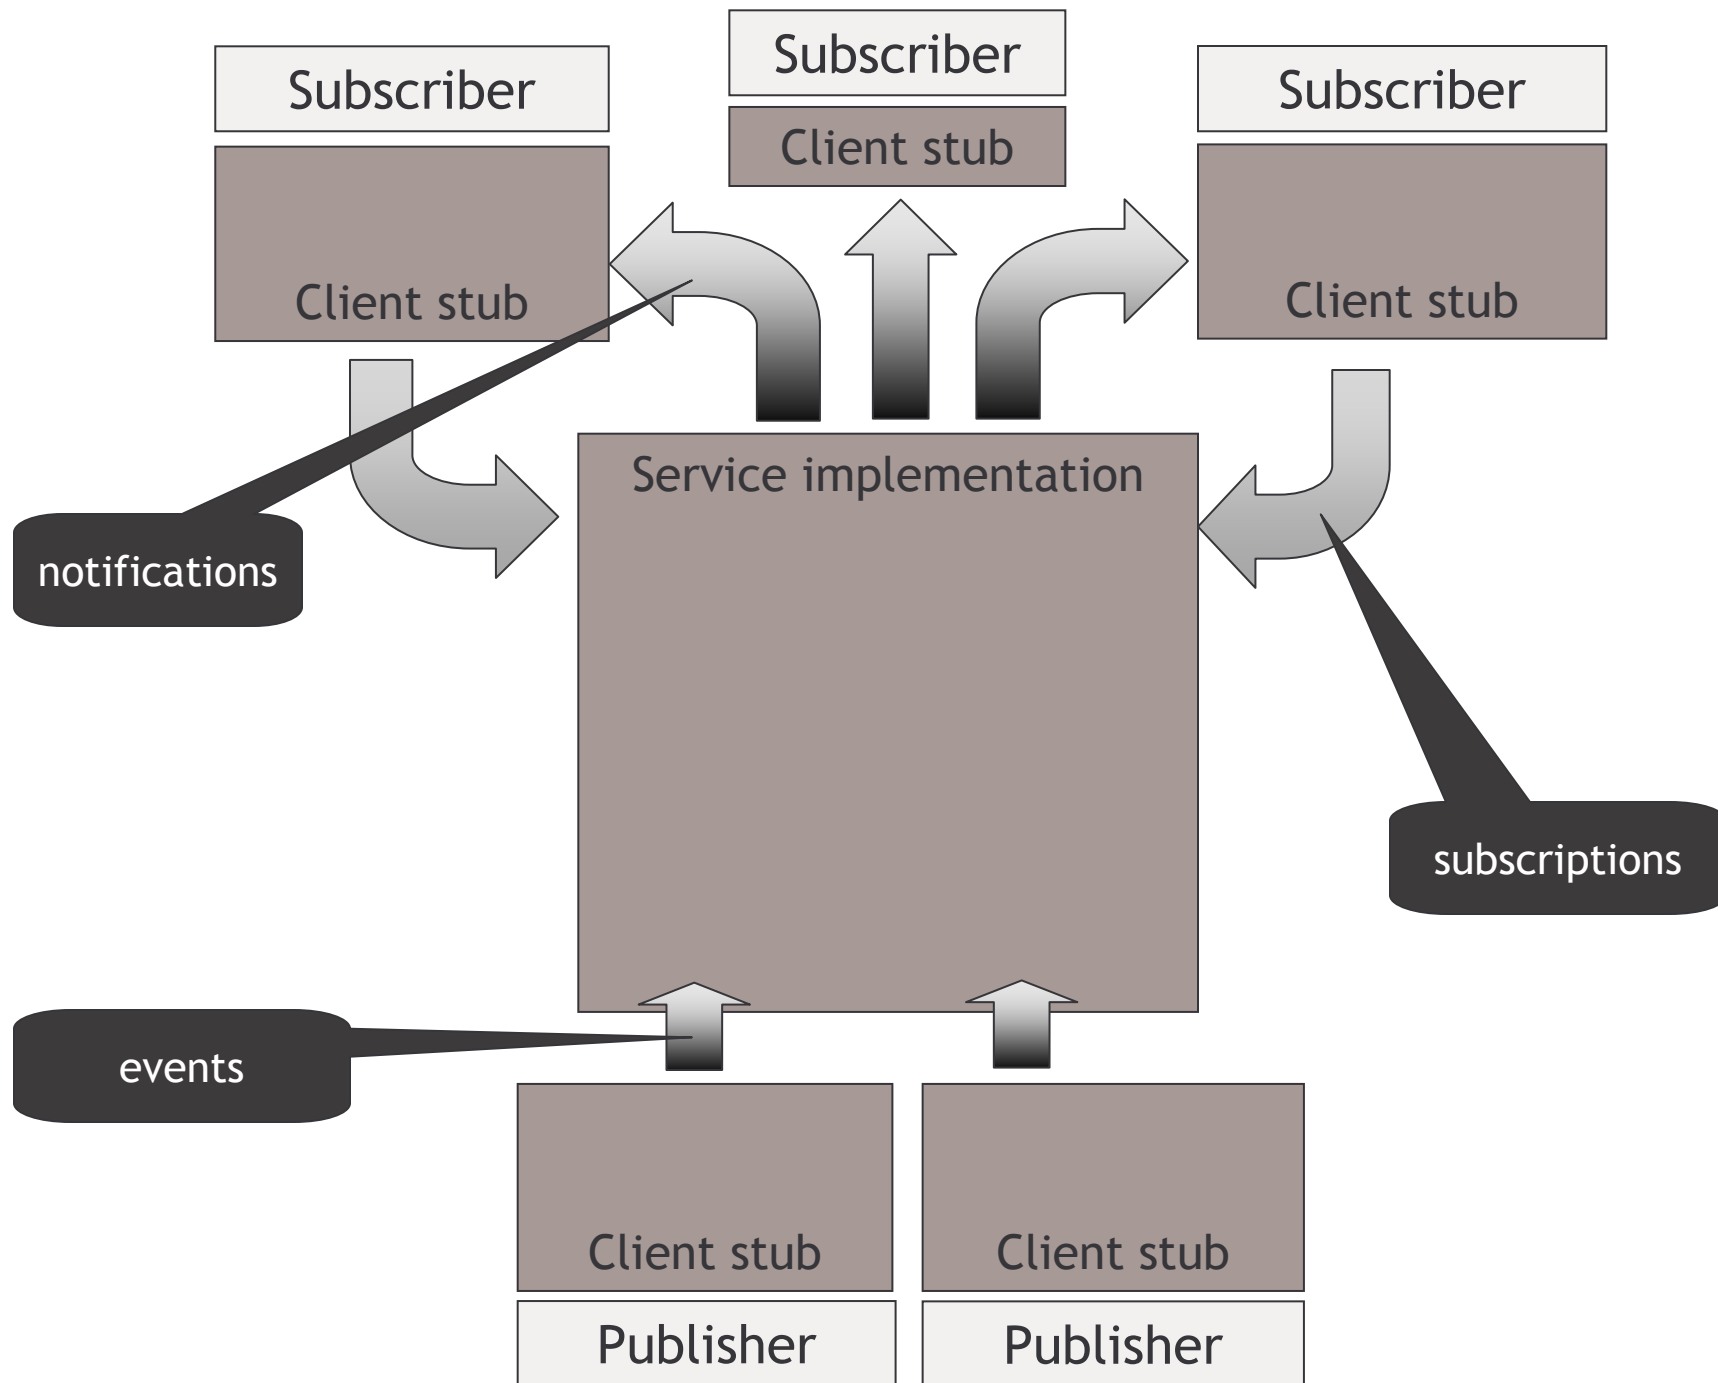
Approach summary

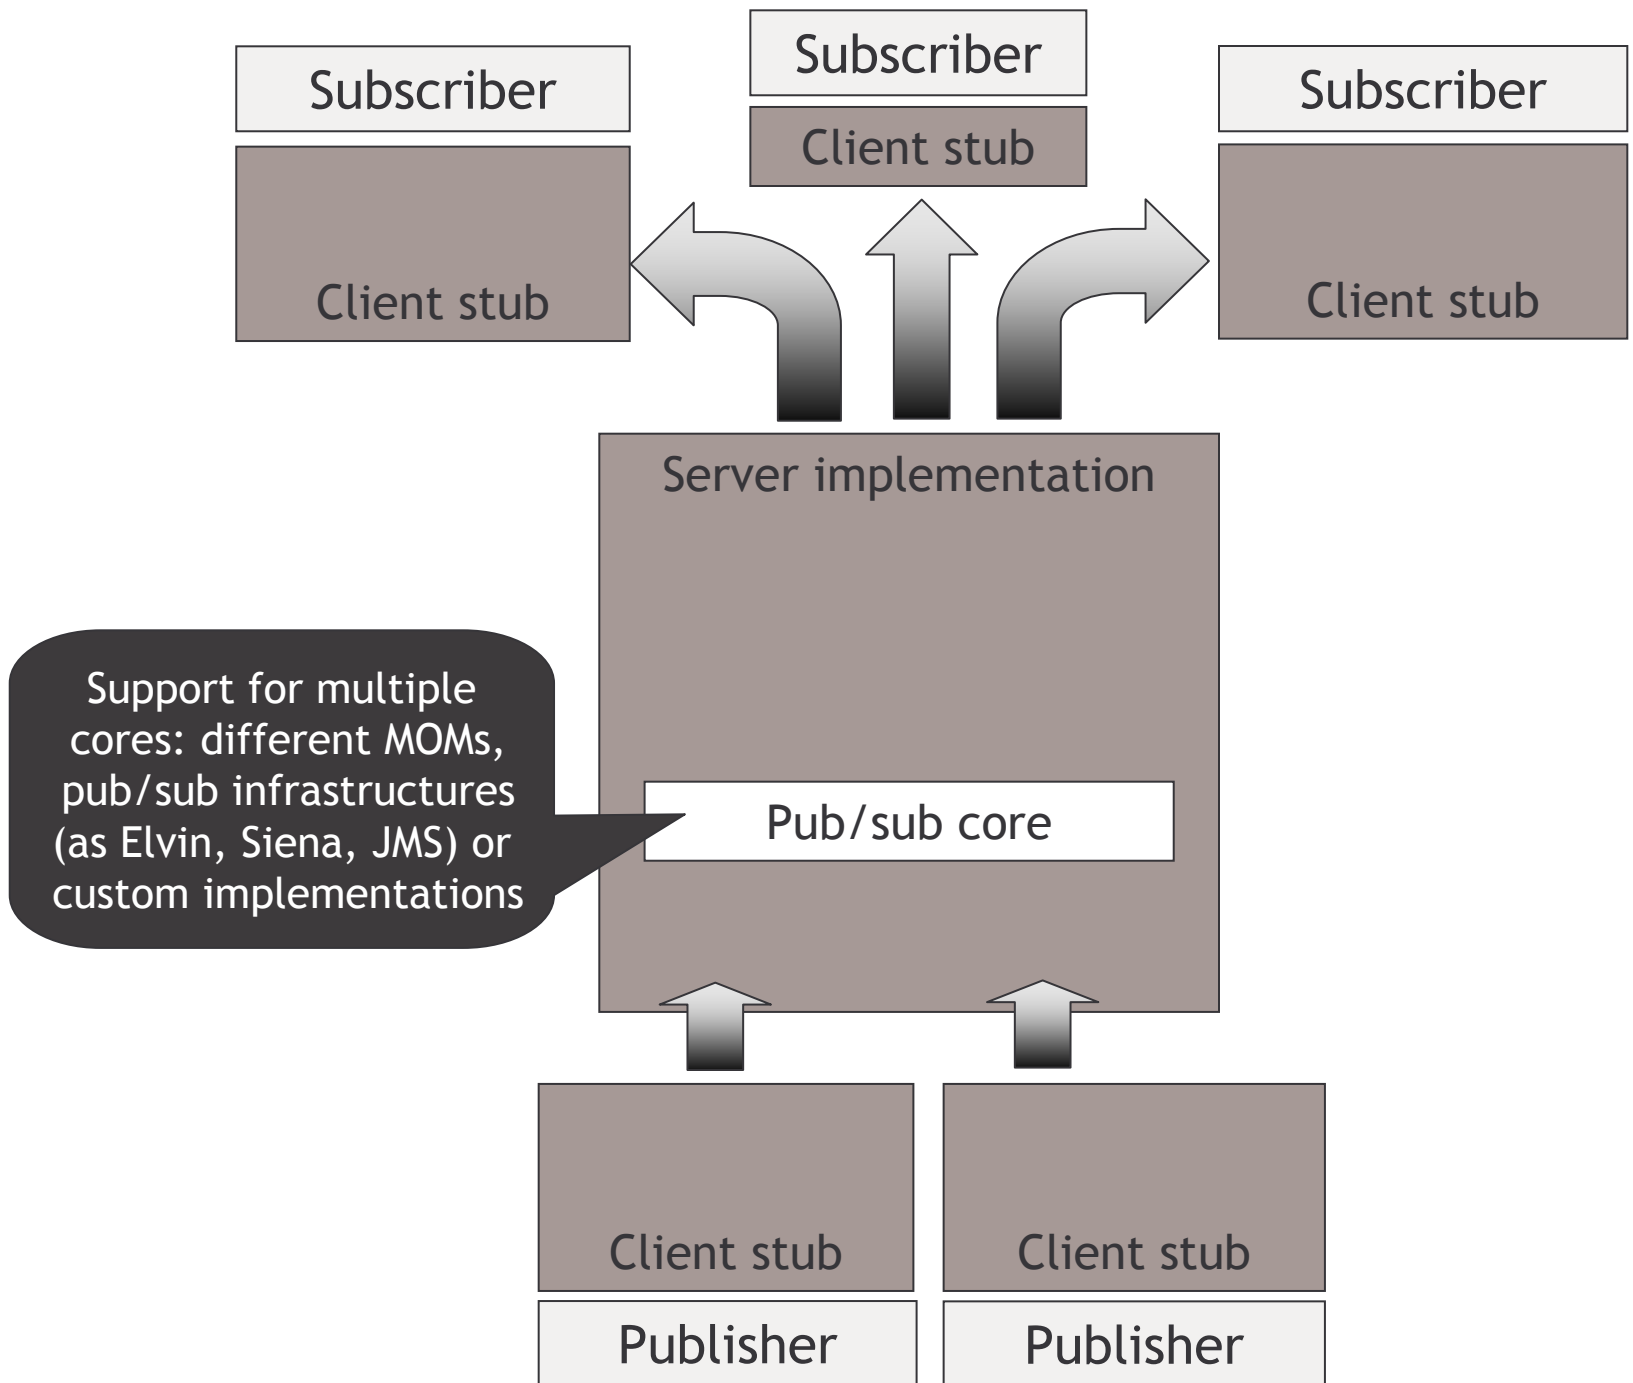


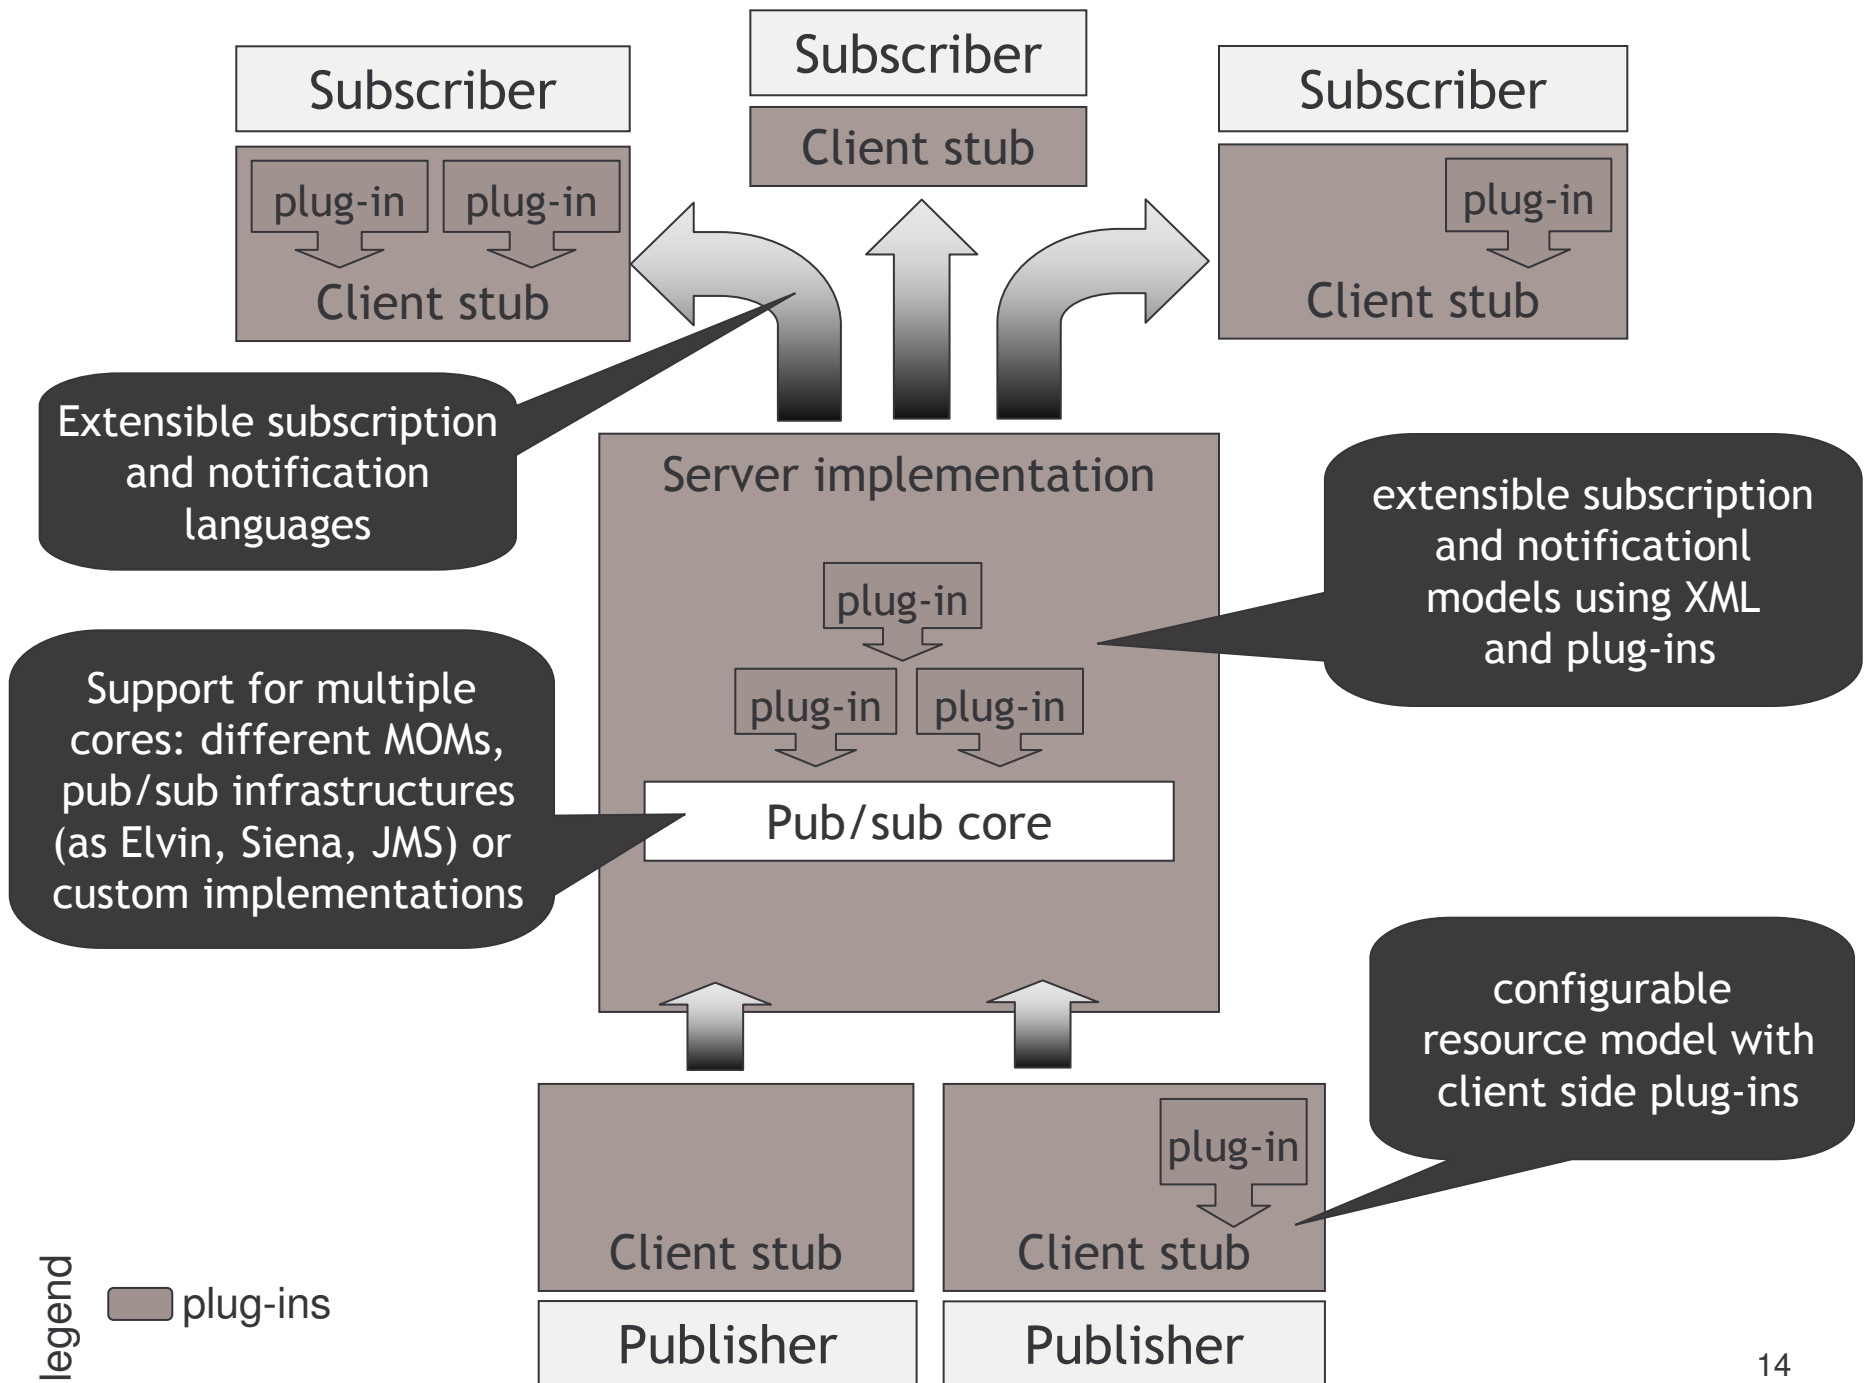
Publish/subscribe design dimensions

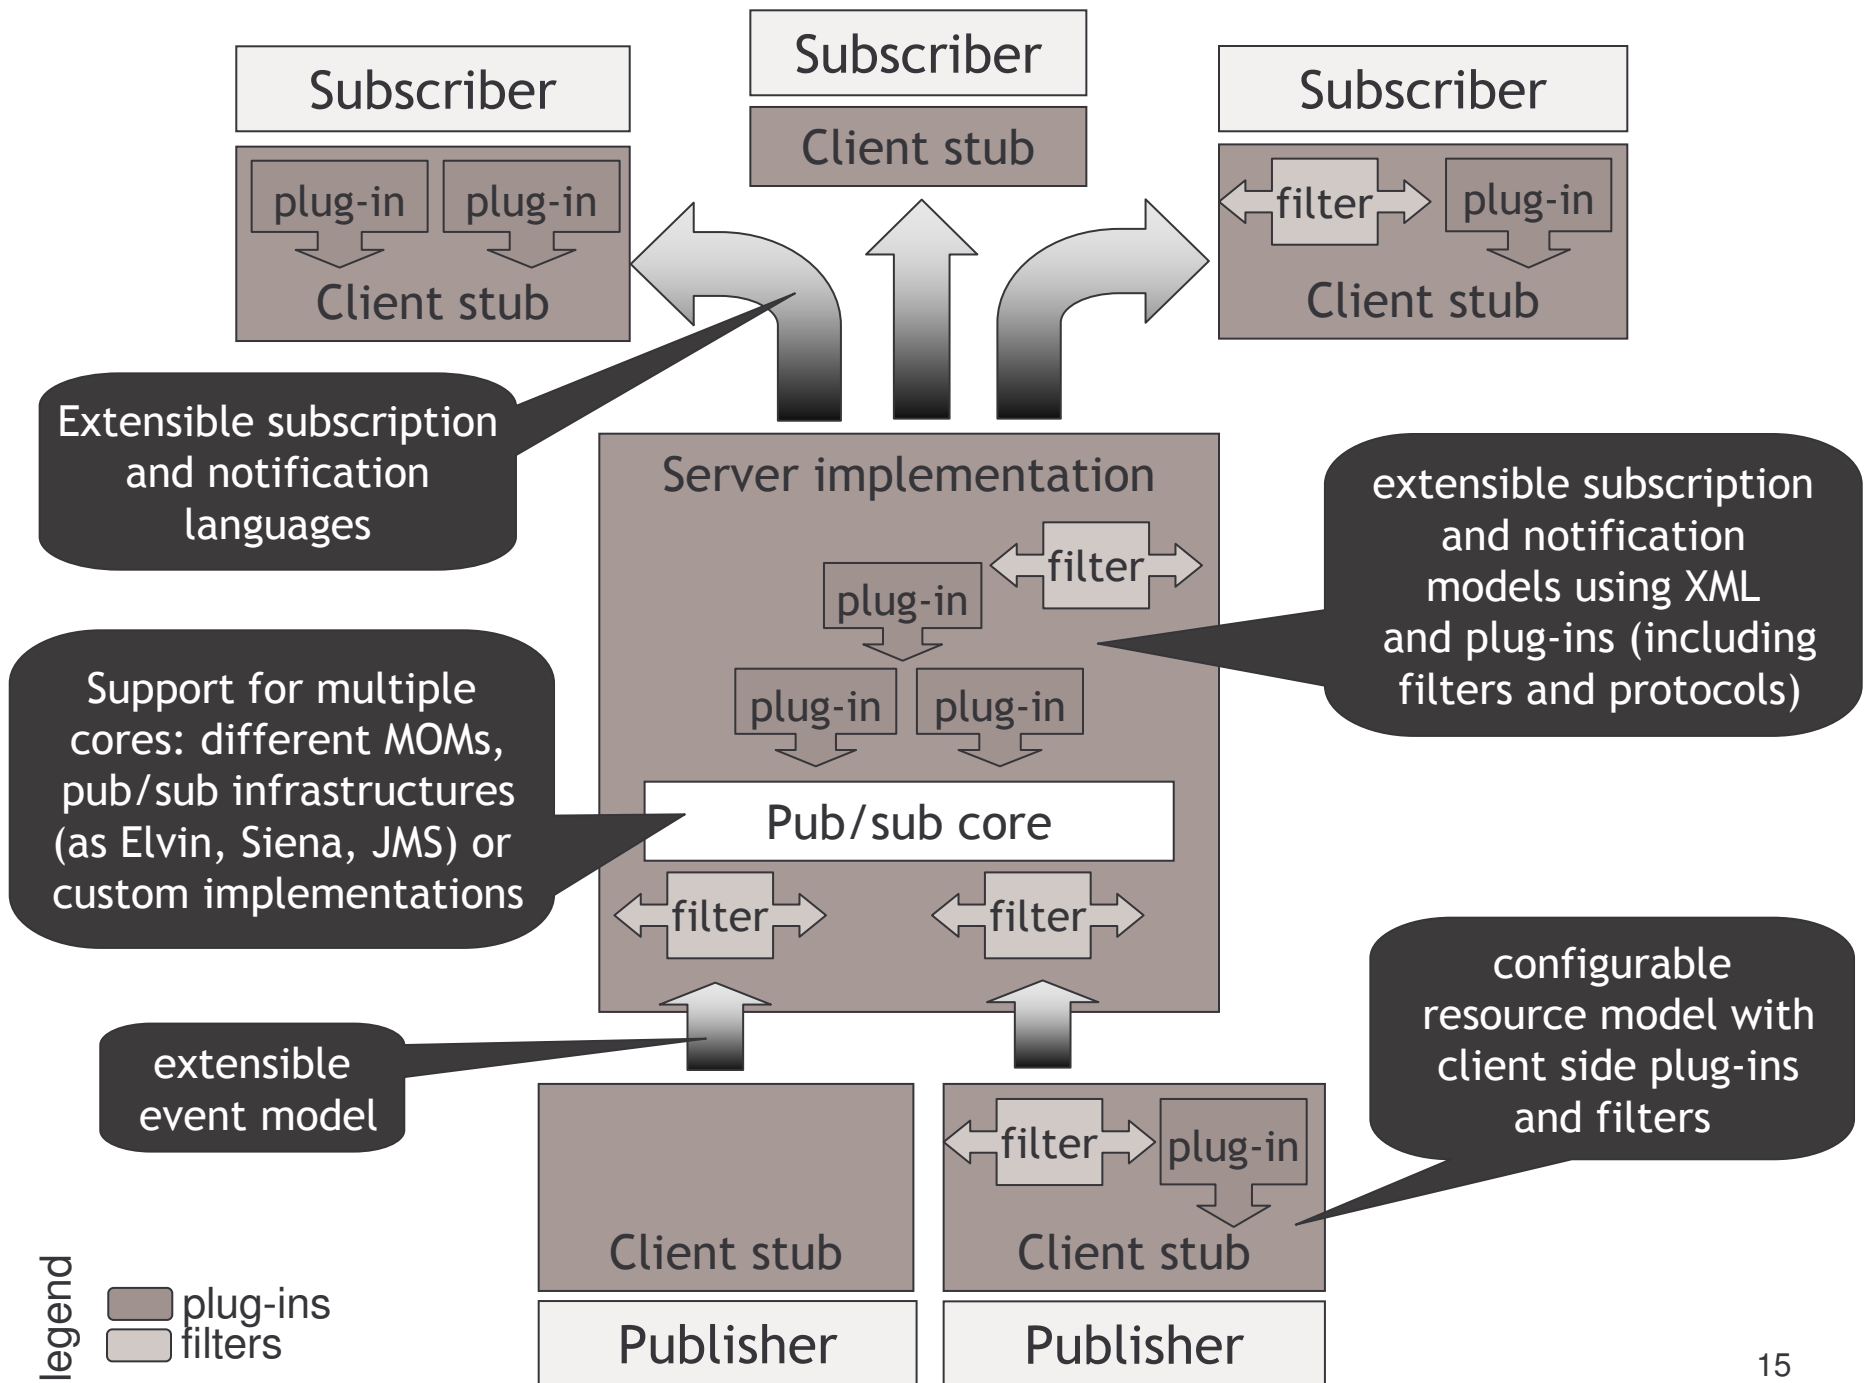
- Extended version (see Protocol*) of Rosenblum and Wolf's model that represent the variability dimensions in our approach

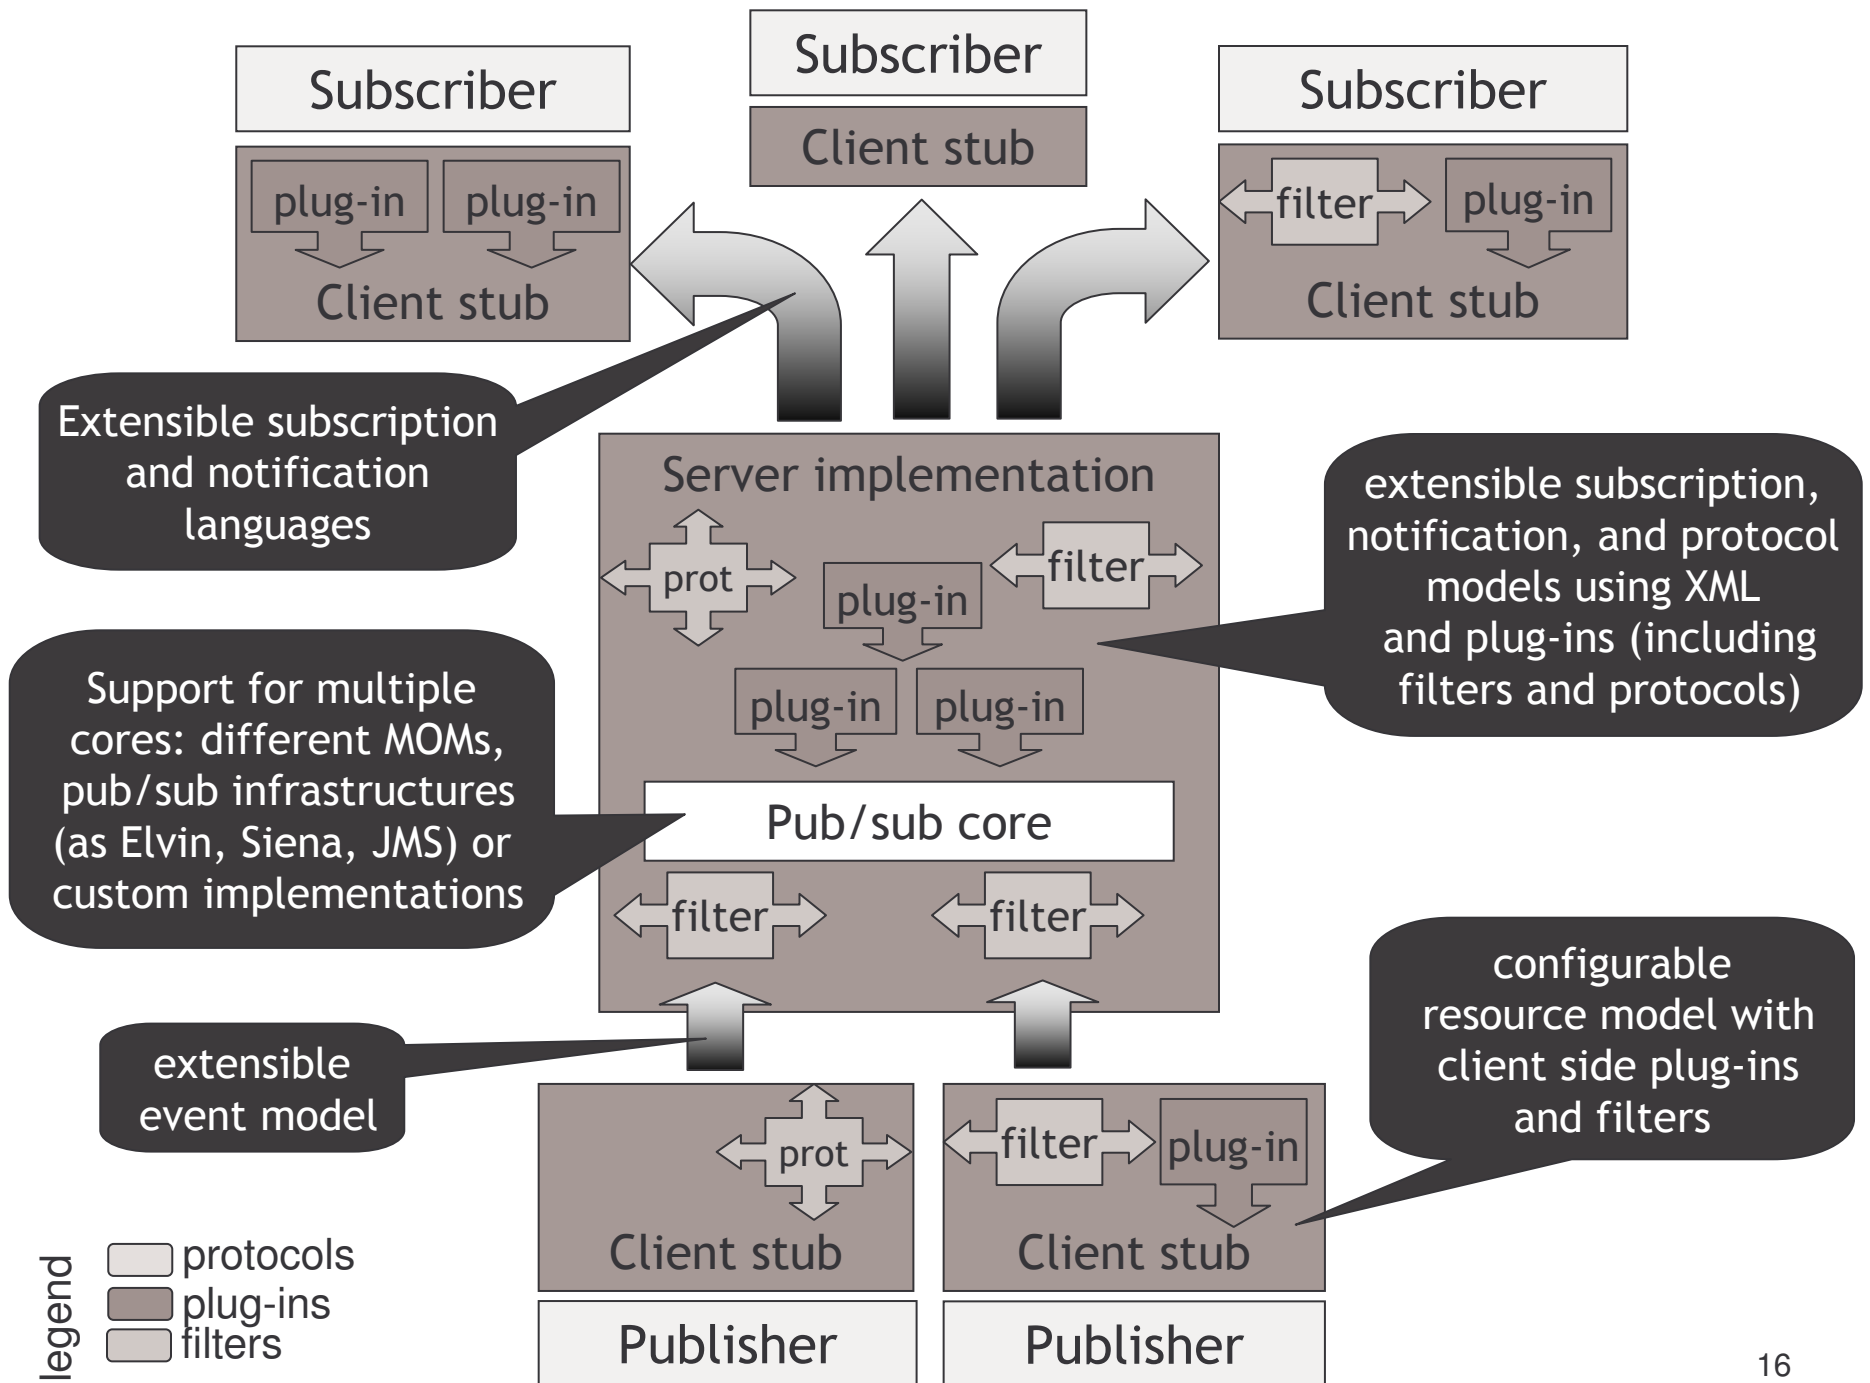
Dimension	Definition	Example
Subscription	specifies how subscribers express interest in subsets of events	content-based, topic-based, advanced event processing
Notification	specifies how notifications are delivered to subscribers	push, pull, both, others.
Event	Specifies how events are represented	tuple-based, record-based, XML documents
Protocol*	other kinds of interaction with the service	Interaction protocols: authentication, manual roaming Infrastructure protocols: federation, replication, fault-tolerance
Resource	defines where in the system (publishers/subscribers/routers) the extensions are placed	client-side, server-side











Variability dimensions summary

Dimension	Approach
Subscription	Extensible subscription language Subscription plug-ins
Notification	Extensible notification language Notification plug-ins
Event	Extensible event representation Filters Event adapters and publish/subscribe cores
Protocol	Protocol plug-ins
Resource	Configuration managers that interpret configuration descriptions Dynamic parsers

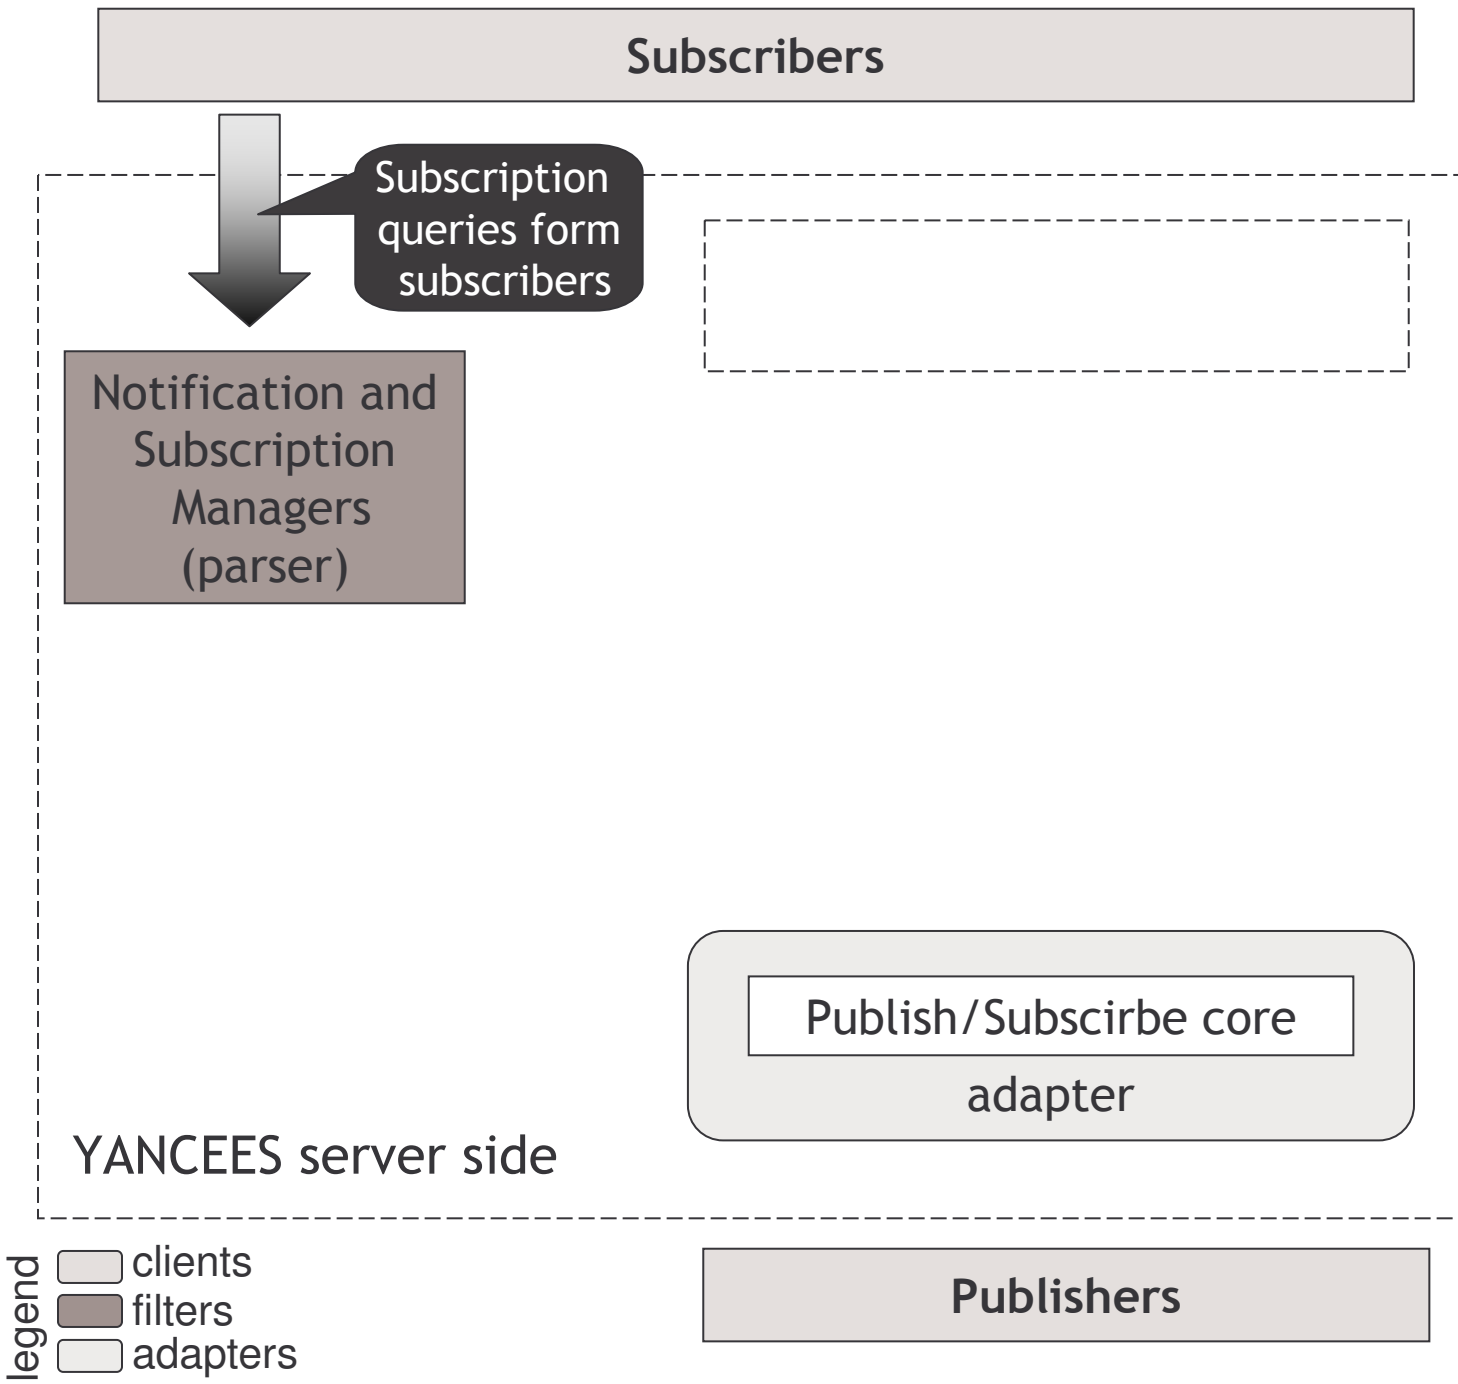
Implementation details

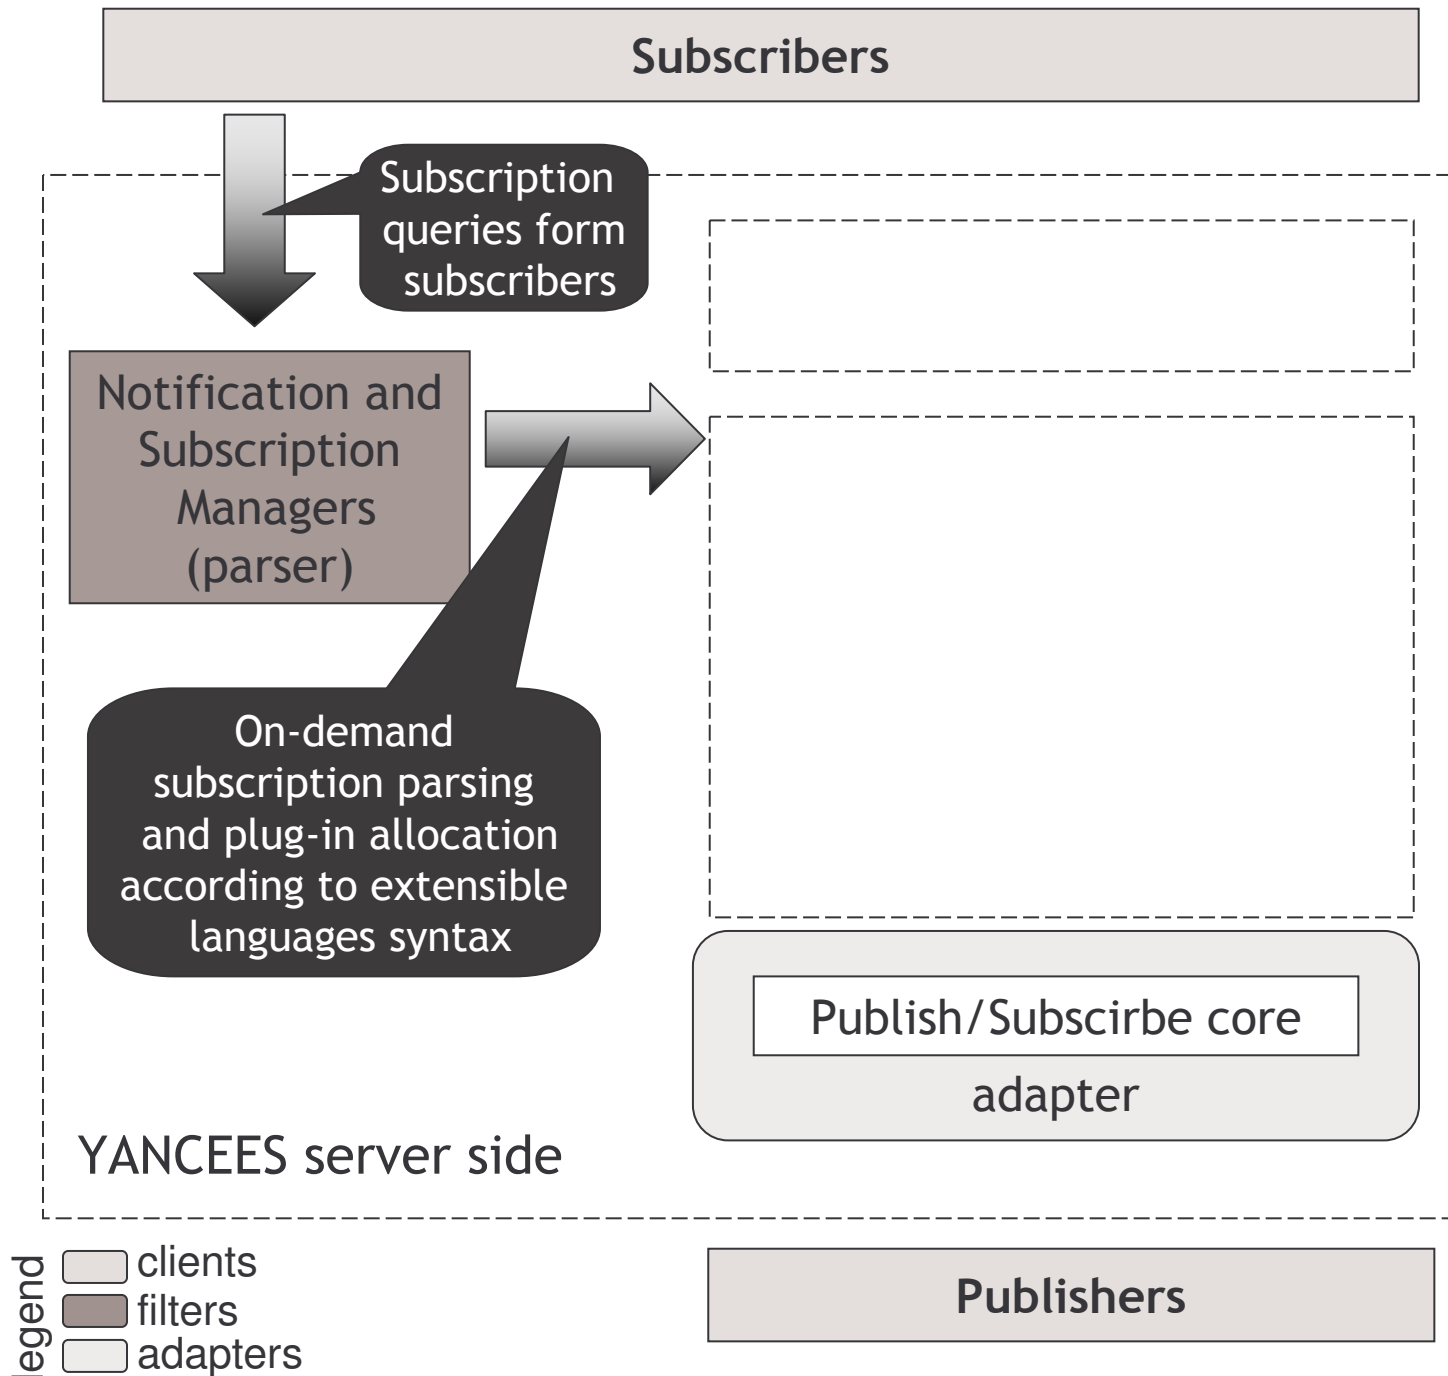
- The application is in Java
- The extensible language used is XML (XMLSchema)
- Events, Subscriptions and Notifications are all represented in XML, as well as the configuration language.
 - Events can also be objects.
- The interaction with the service (pub/sub API) is done through RMI
- Protocol plug-in interfaces are currently using RMI
- Siena, Elvin and a custom topic-based switcher were used as the basic pub/sub cores

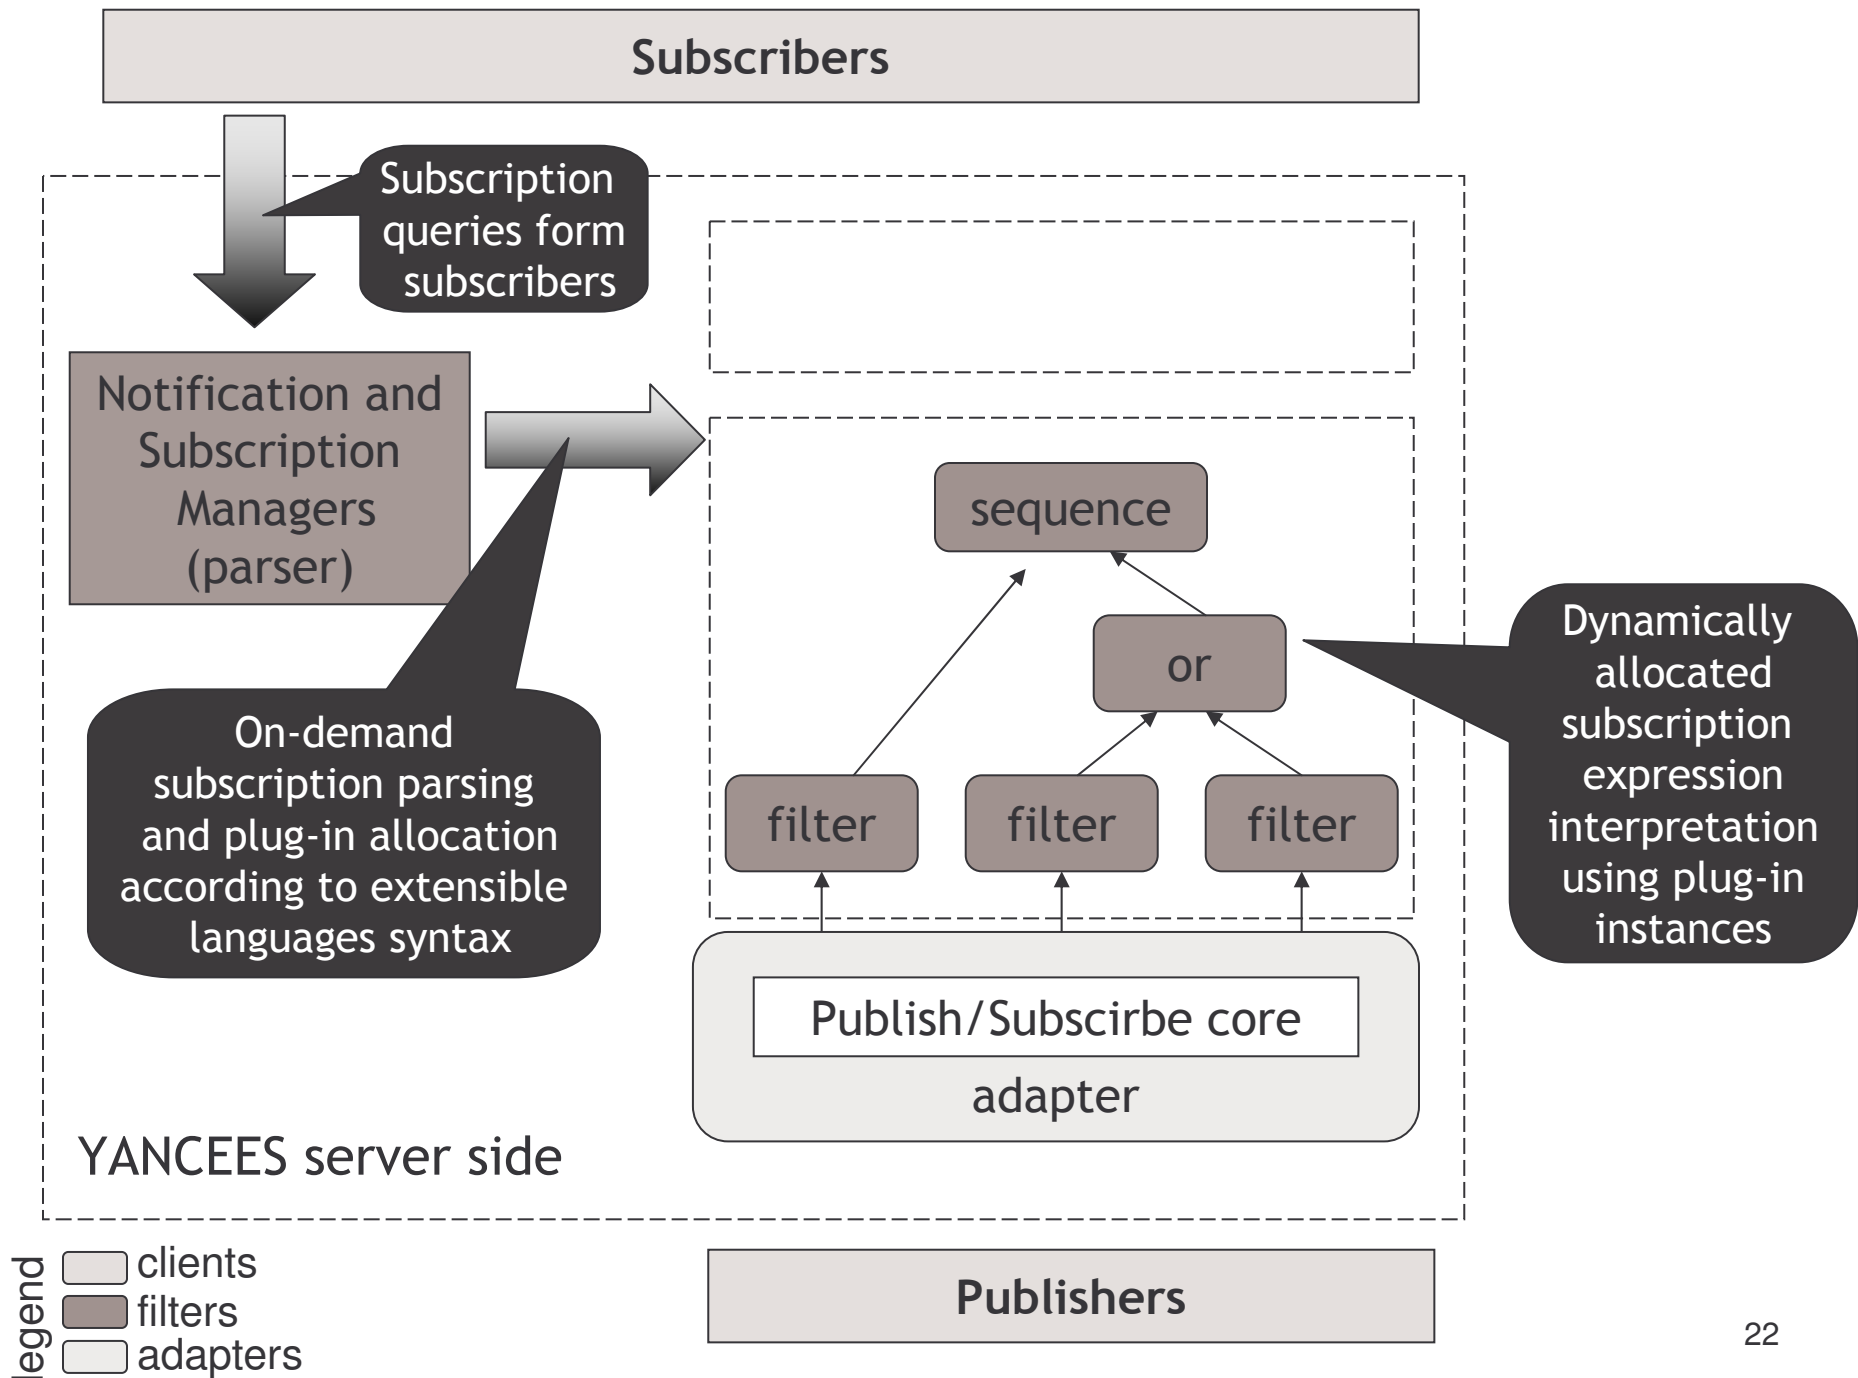
Example of dynamic parsing of subscriptions

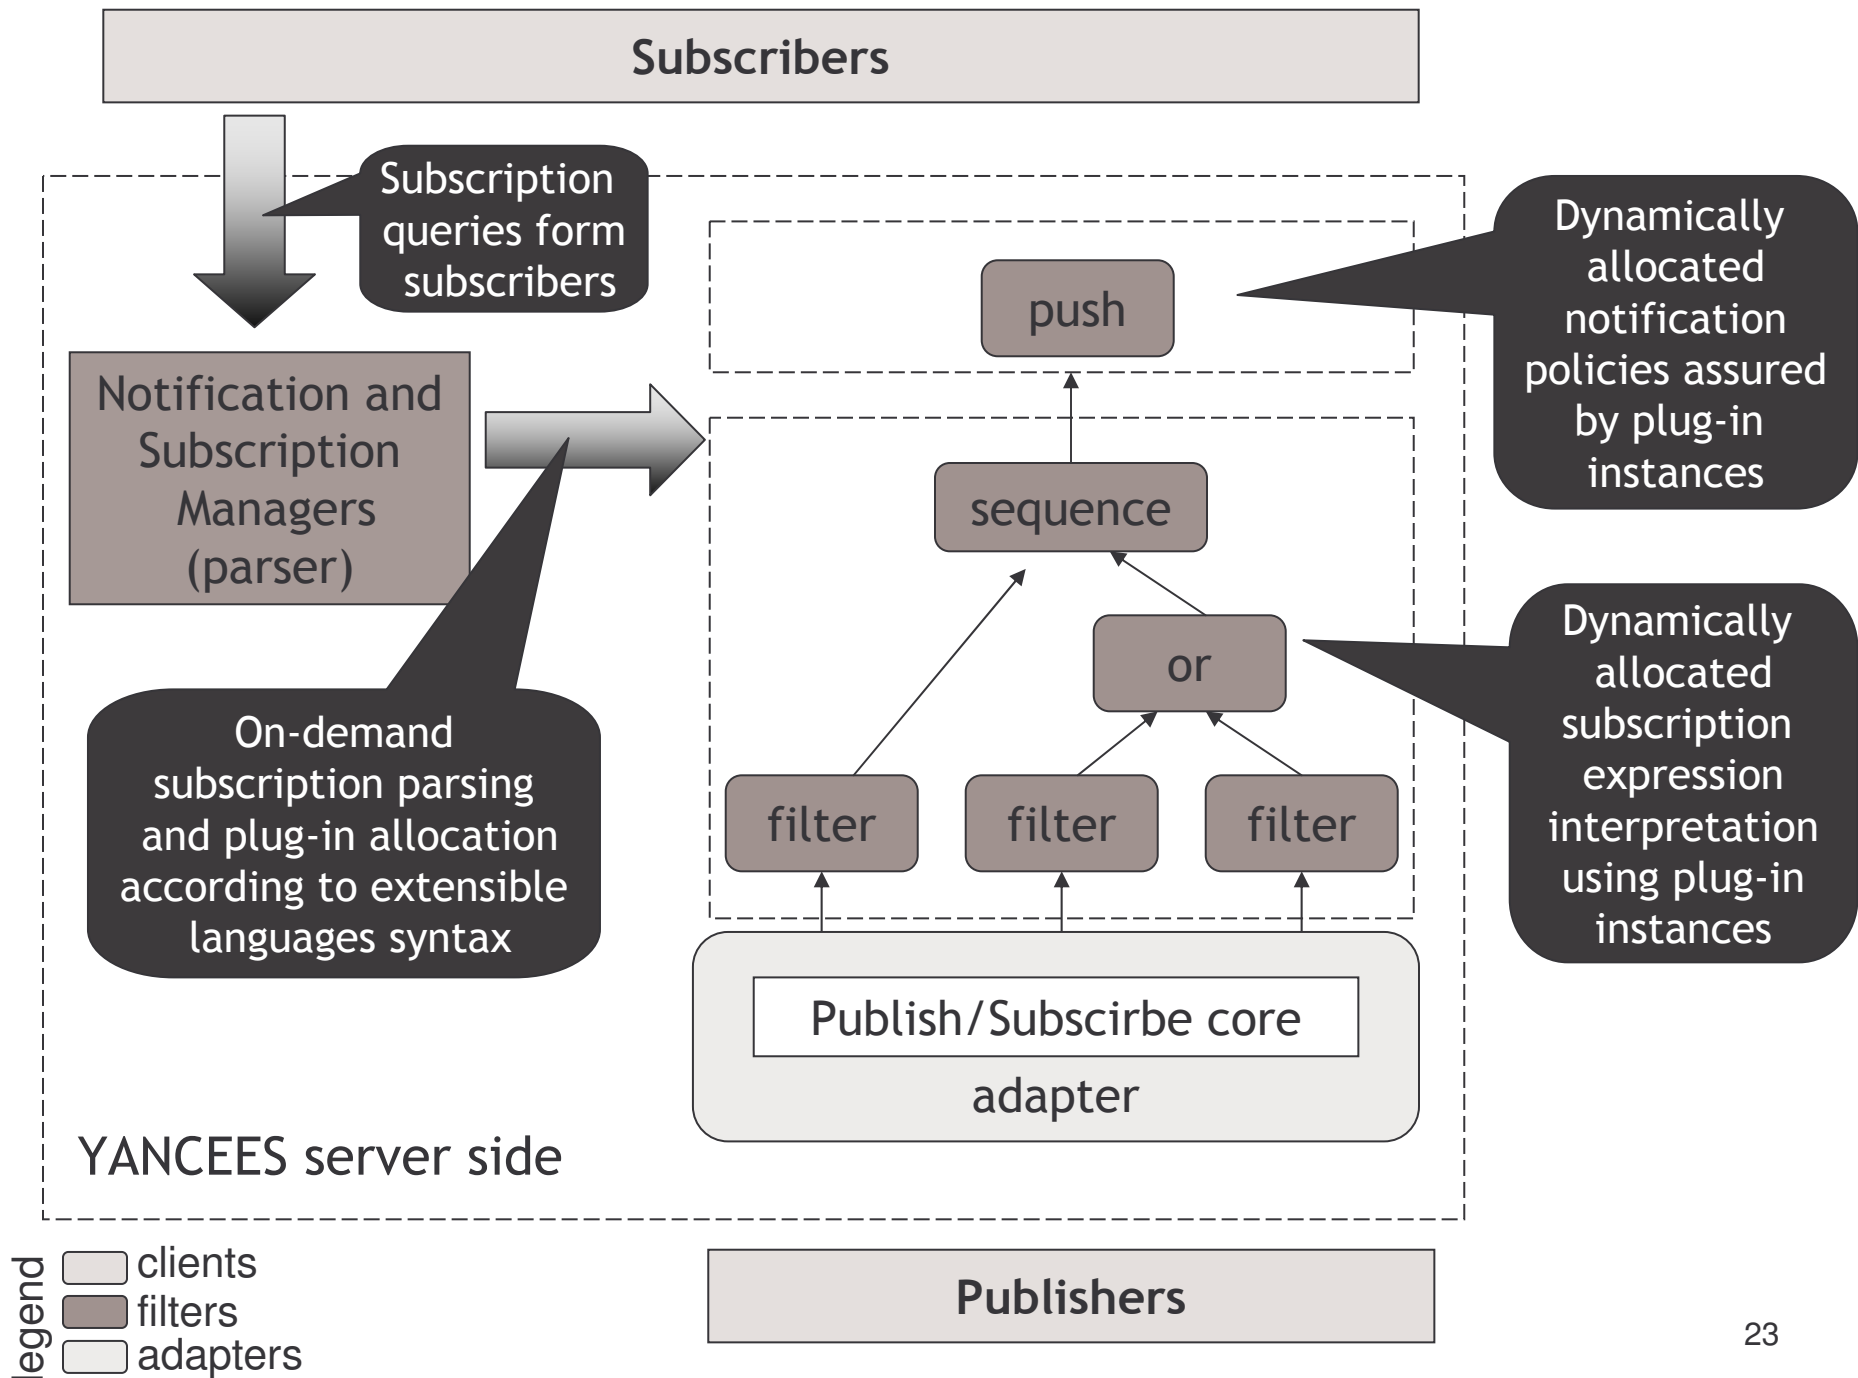
- Consider the following sequence detection subscription as an example:

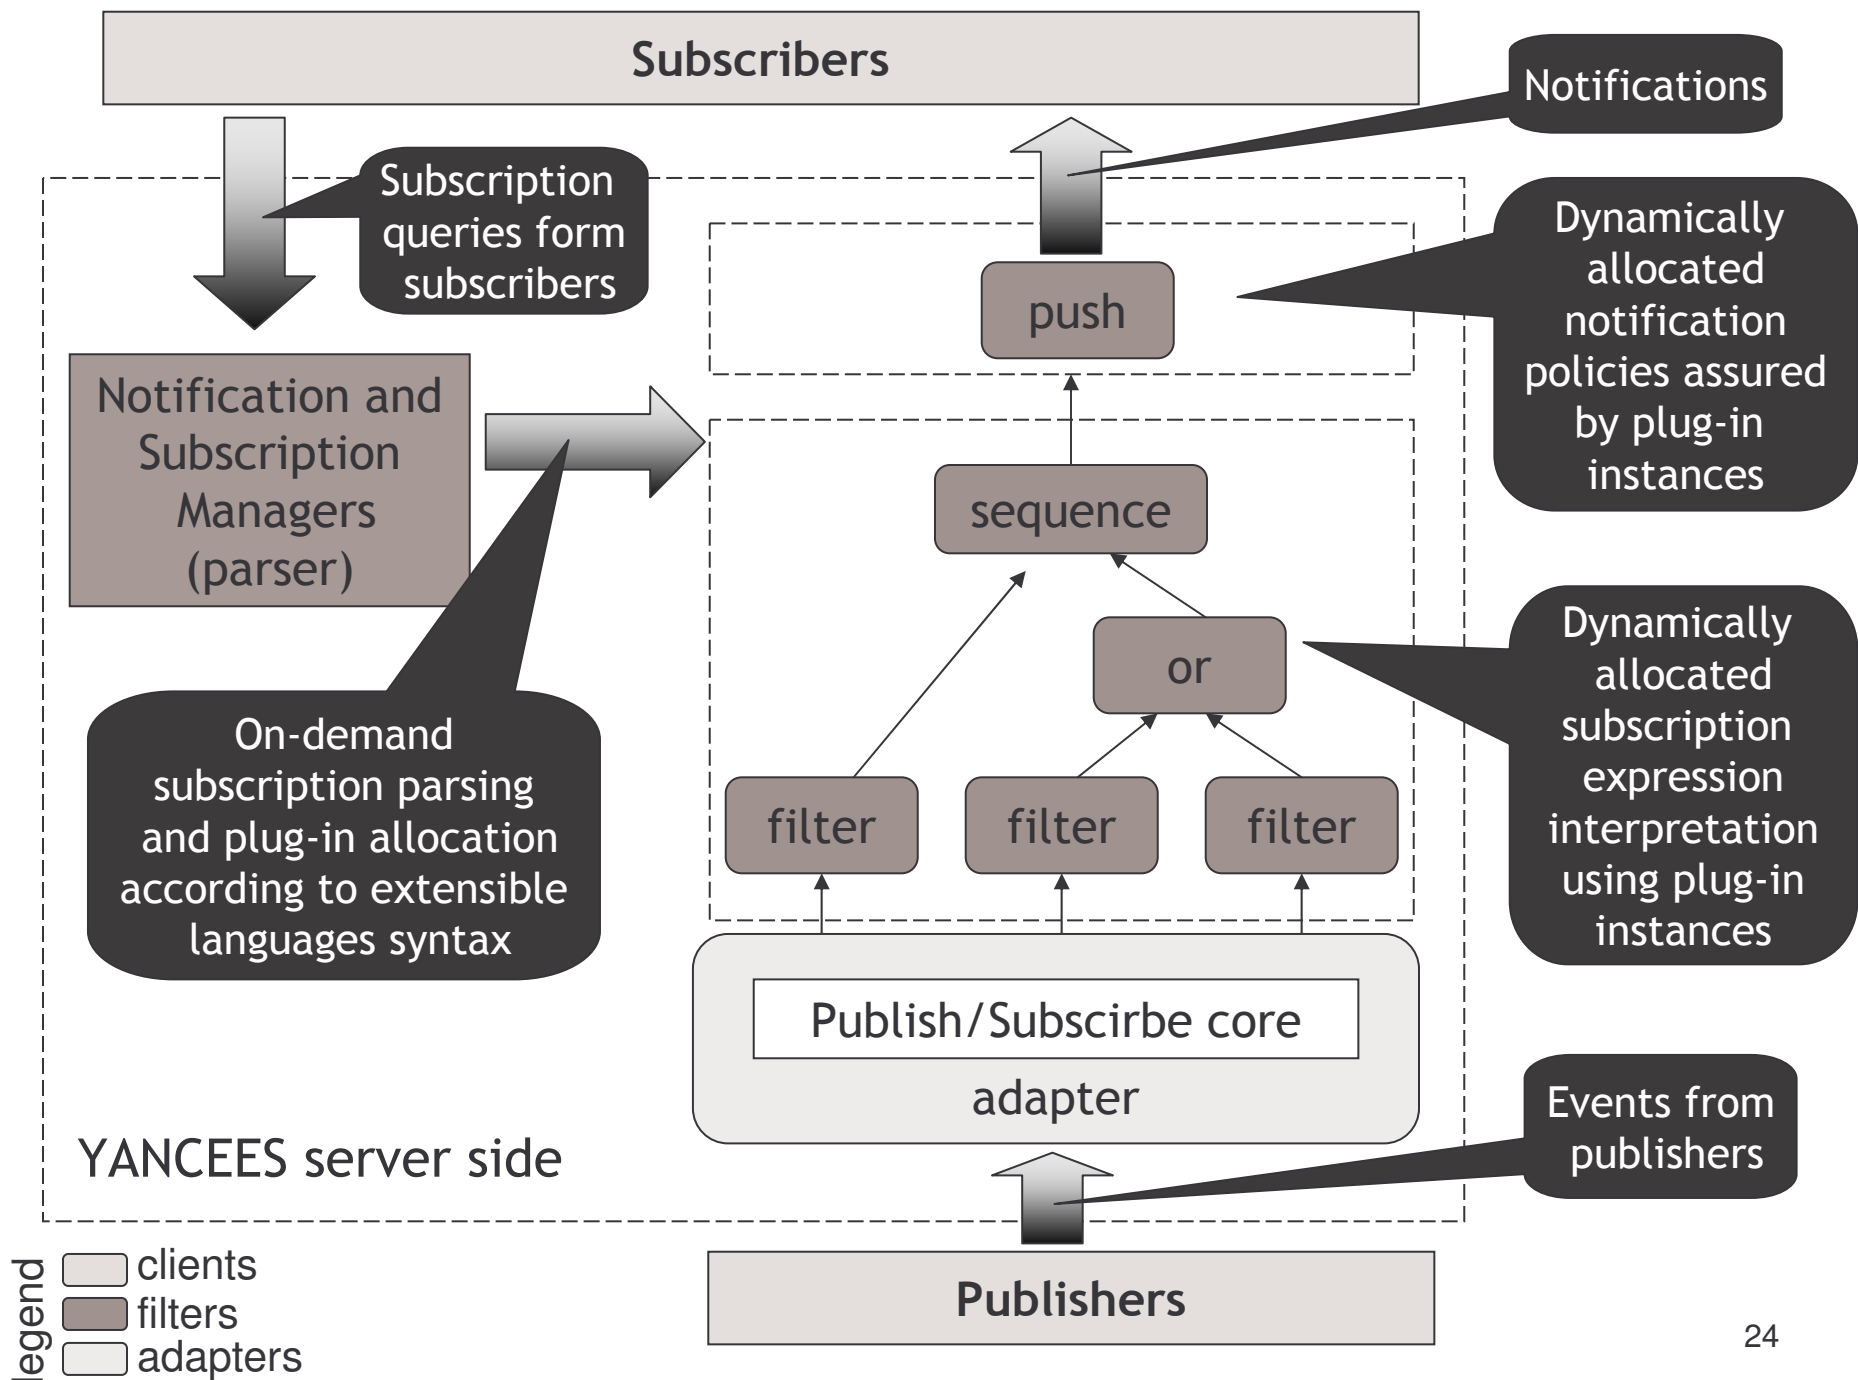
```
<subscribe>
  <sequence>
    <or>
      <filter>
        <eq> <name> name </name>
          <value> Roberto </name> </eq>
      </filter>
      <filter> ... </filter>
    </or>
    <filter> ... </filter>
  </sequence>
  <push>
</subscribe>
```



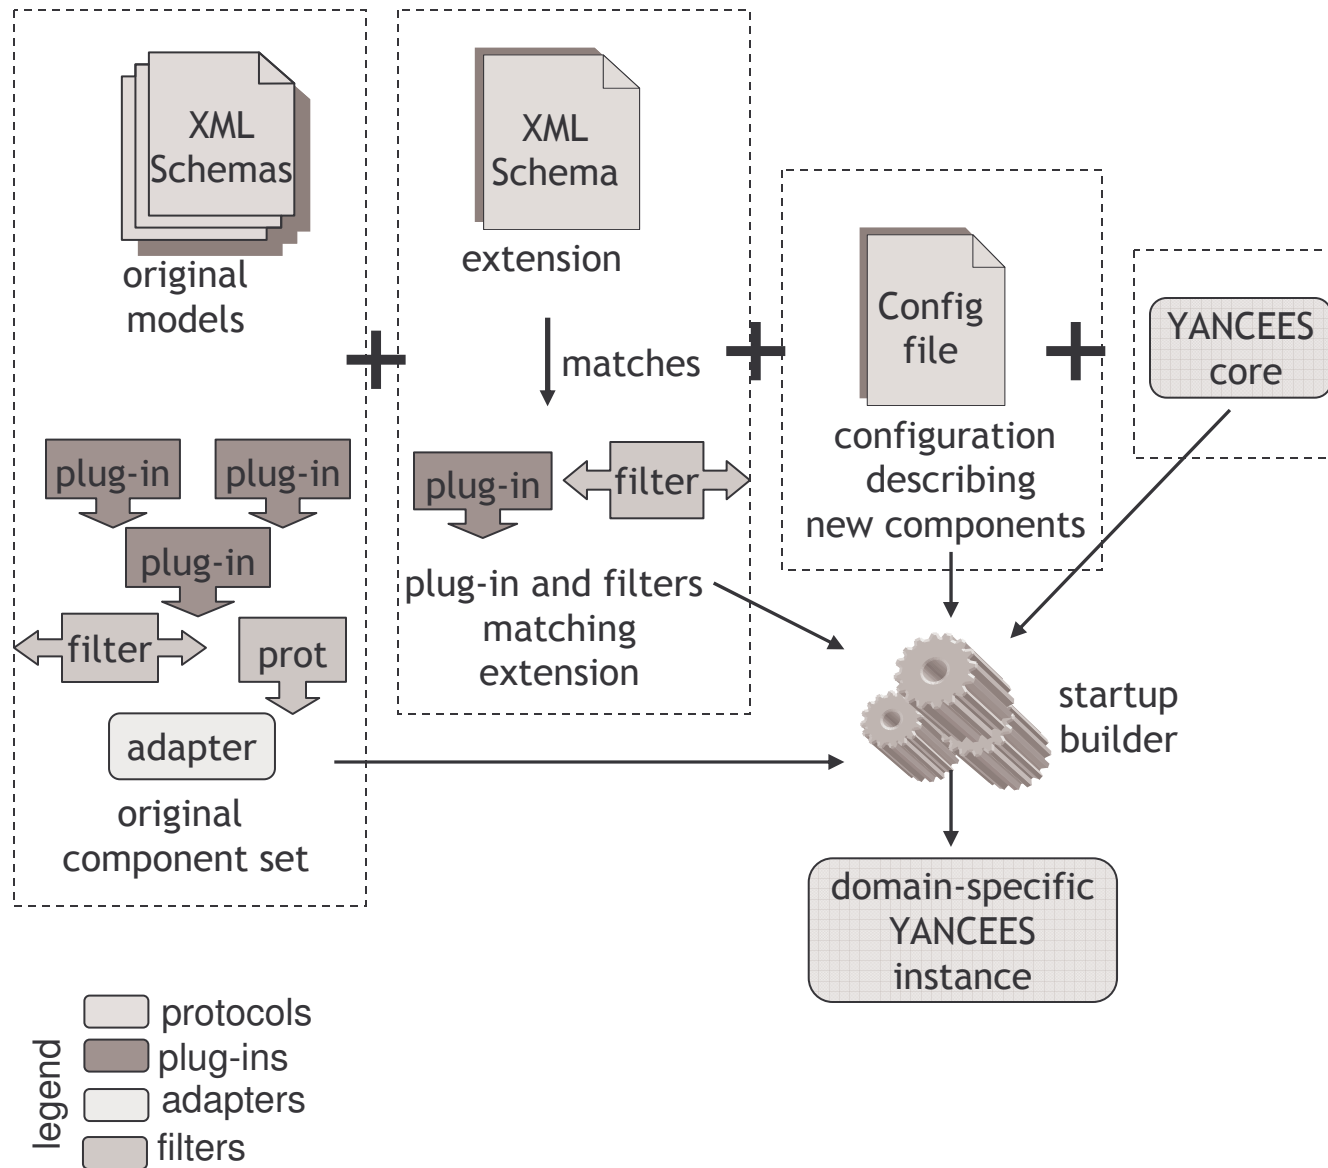








How to extend YANCEES?



1. Define a language extension using XML Schema
2. Implement plug-ins and filters for that extension
3. register plug-in in the configuration file
4. Restart the server with the new configuration

How does it compare to other approaches? (e.g. Elvin, Siena or CORBA-NS)

- None of them are programmable. And they are not easily extensible.
- Siena and Elvin, for e.g., provide content-based routing and sequence detection with push notification only
- CORBA-NS allows the selection of notification (push, pull) and subscription policies to use (channel, topic). The event model is fixed (object-based) at runtime.
 - It is monolithic and no additional features can not be easily added to it or removed from it.

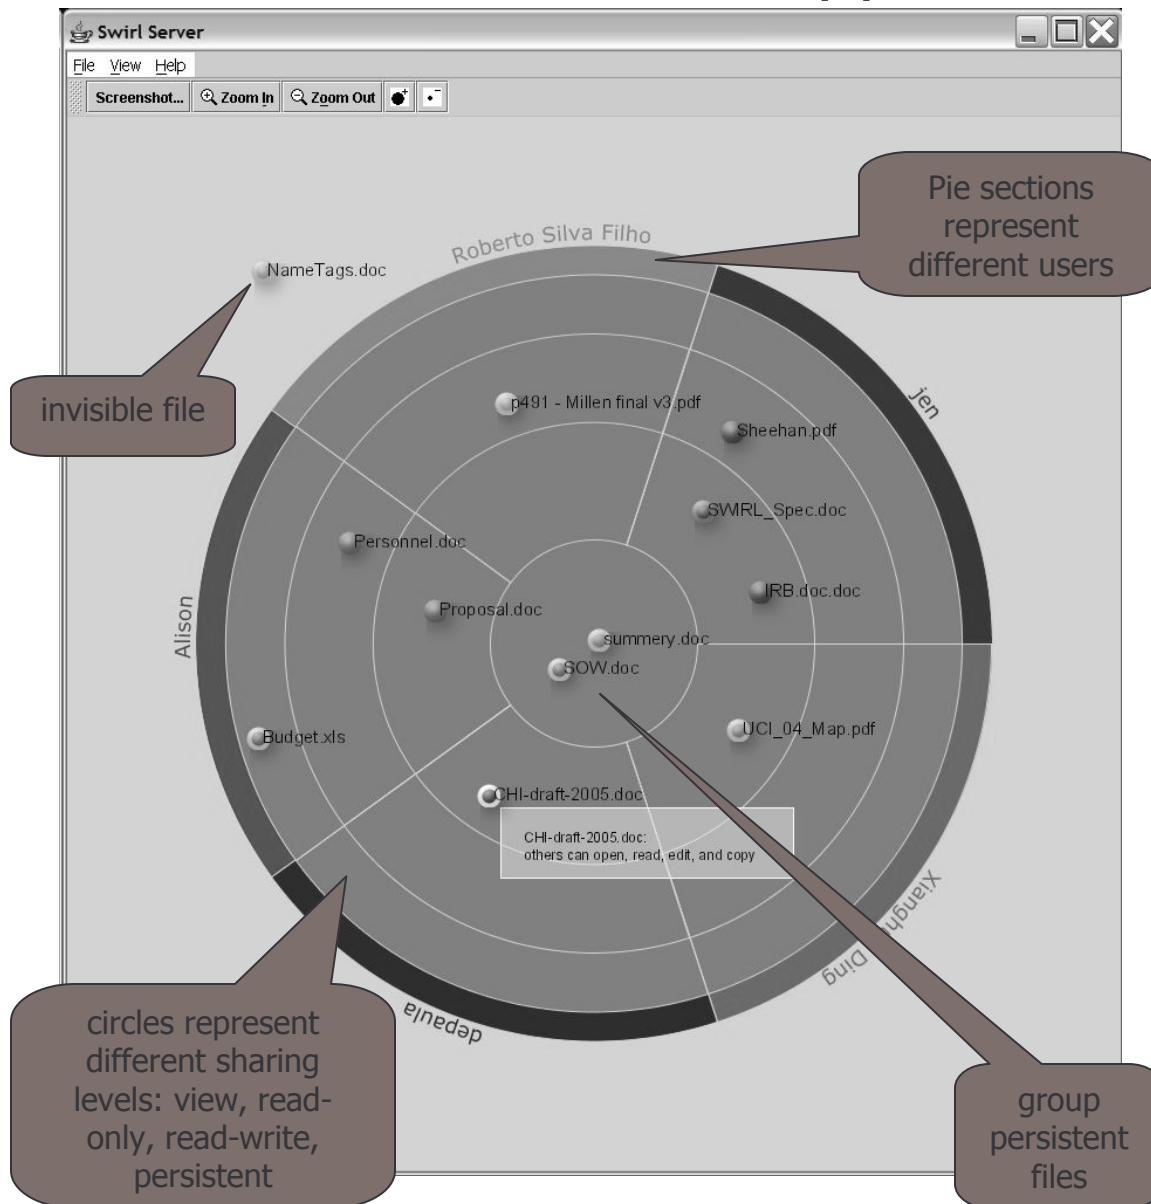
Three example applications and their configurations

- Implementation of a peer-to-peer event bus for ad-hoc file sharing application
- support for a software visualization tool and network activity monitoring application
- Implementation of an awareness server (CASSIUS equivalent)

Peer-to-peer file sharing tool

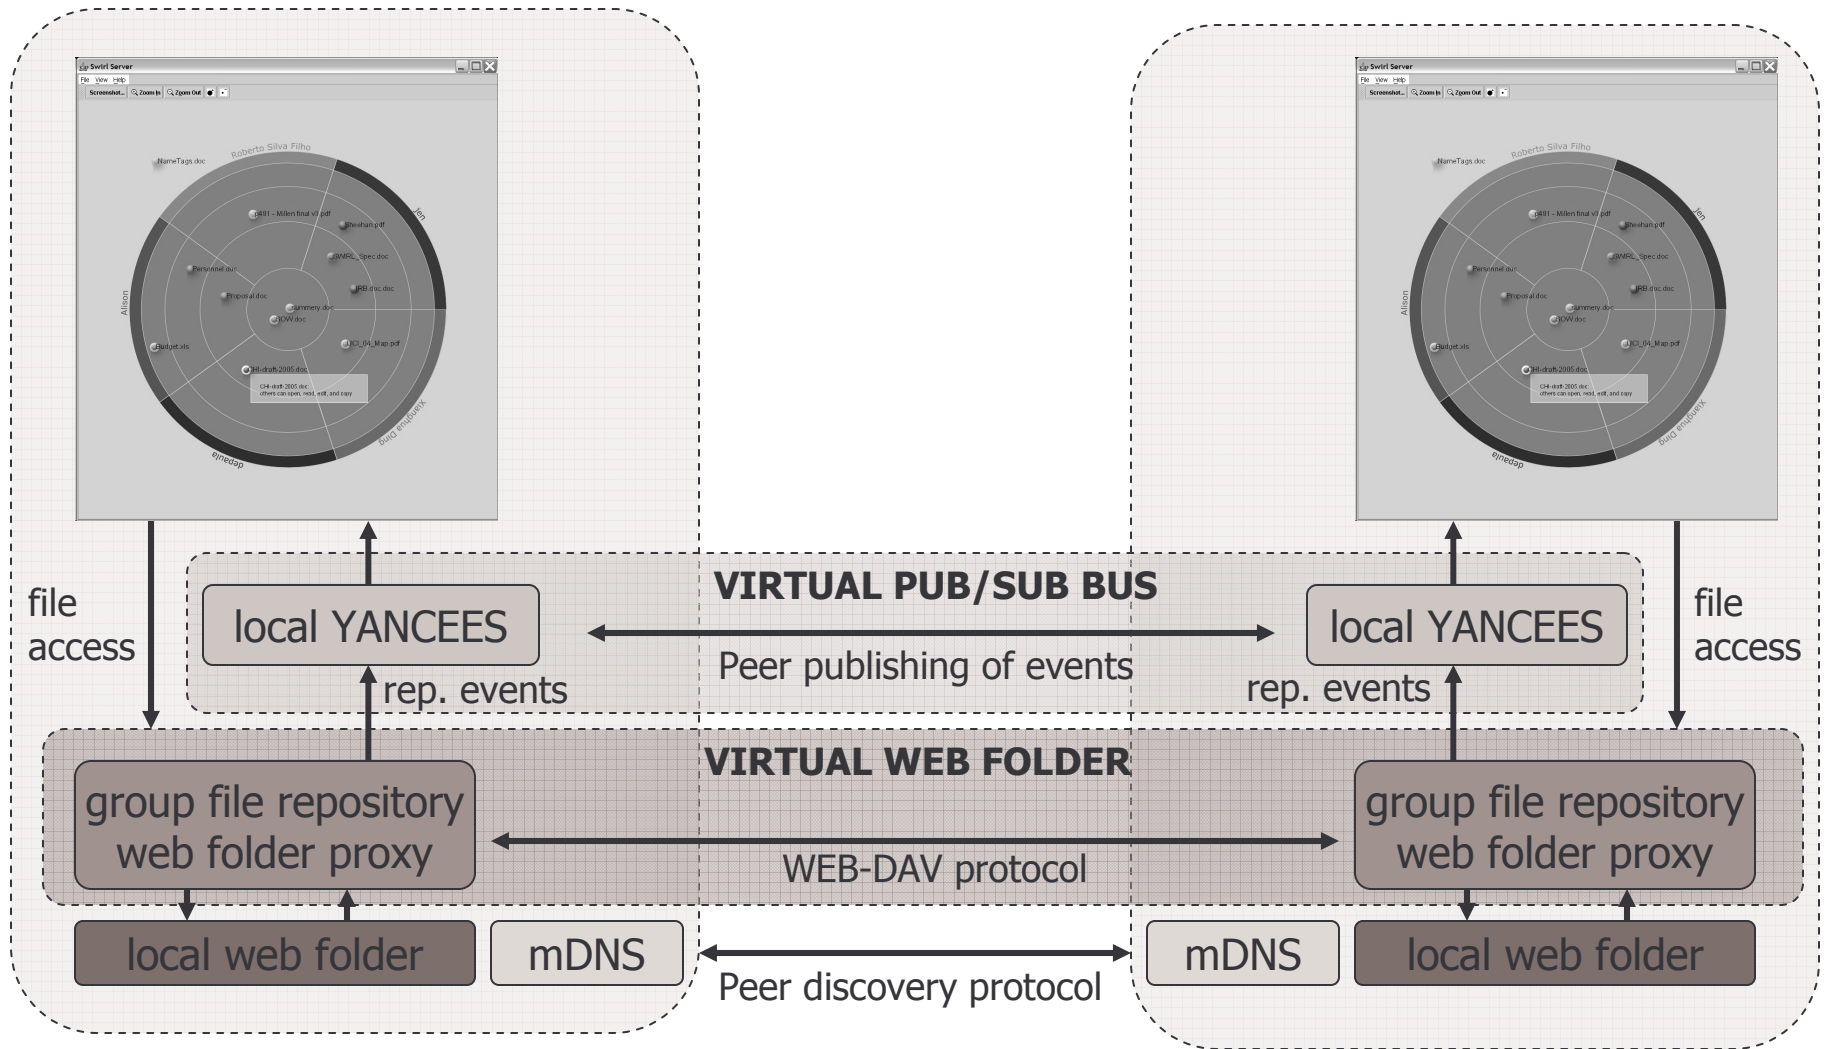
- YANCEES is used to provide a P2P event bus that supports:
 - dynamic peer discovery
 - peer-to-peer publishing

Impromptu: a peer-to-peer, ad-hoc, file sharing application



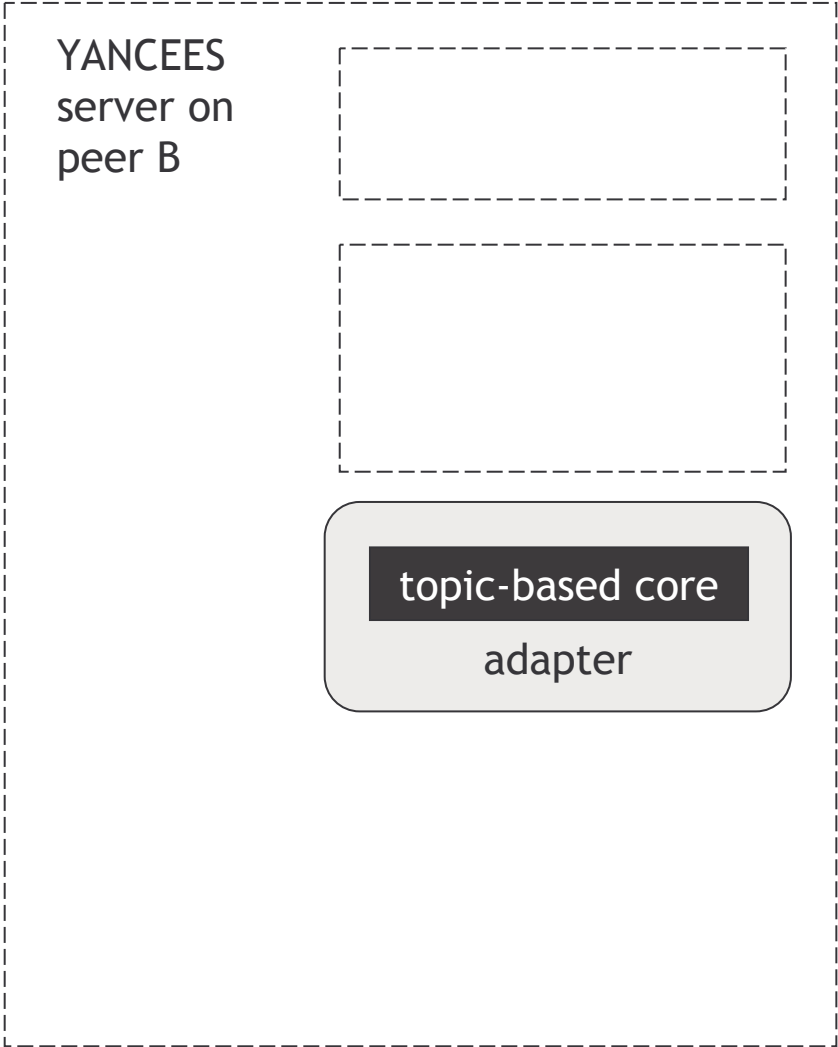
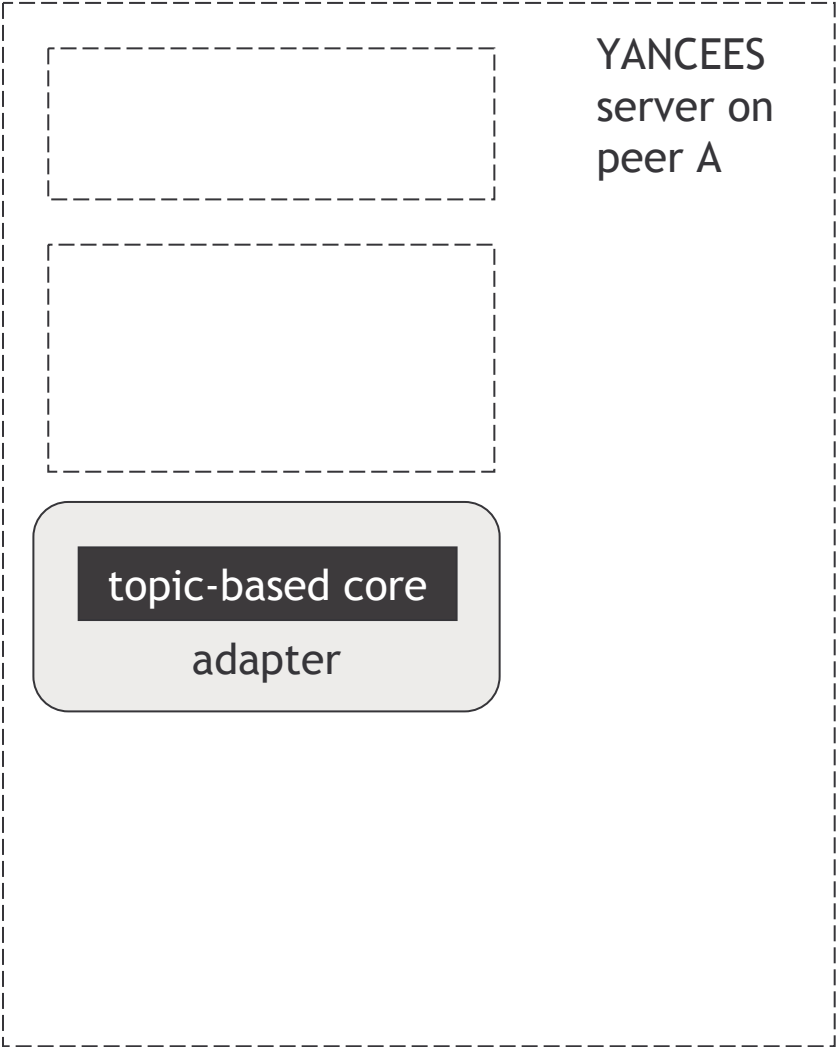
- Different pie slices represent peers in the network
- Users can drag and drop files to be shared
- Concentric circles define different sharing levels
- Files placed in the center are persistent and available to all the members of the group
- Files outside the pie are not shared and become invisible to other peers
- Files blink with the peer color whenever someone reads or modifies it

Impromptu architecture: peer-to-peer file sharing tool support




Subscribers

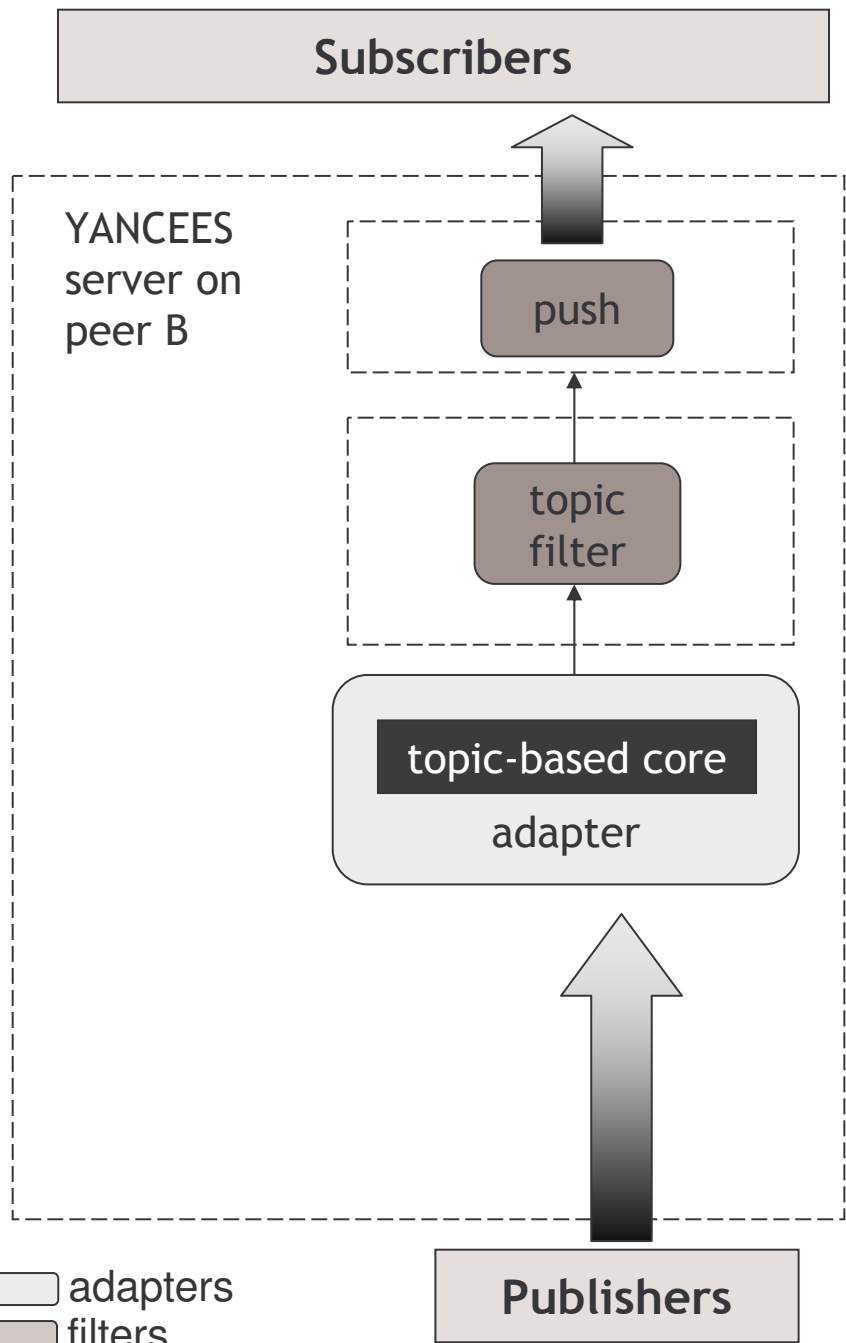
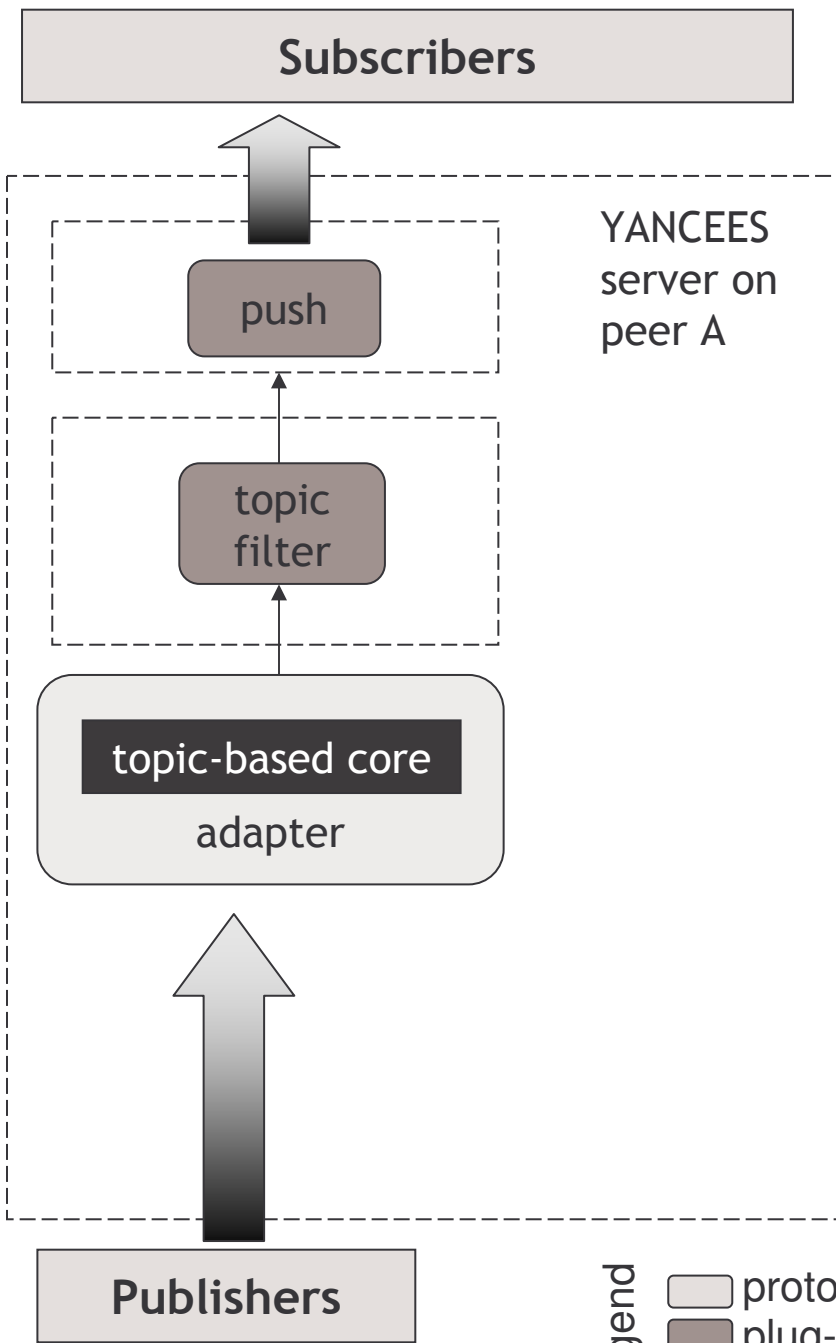
Subscribers







Publishers

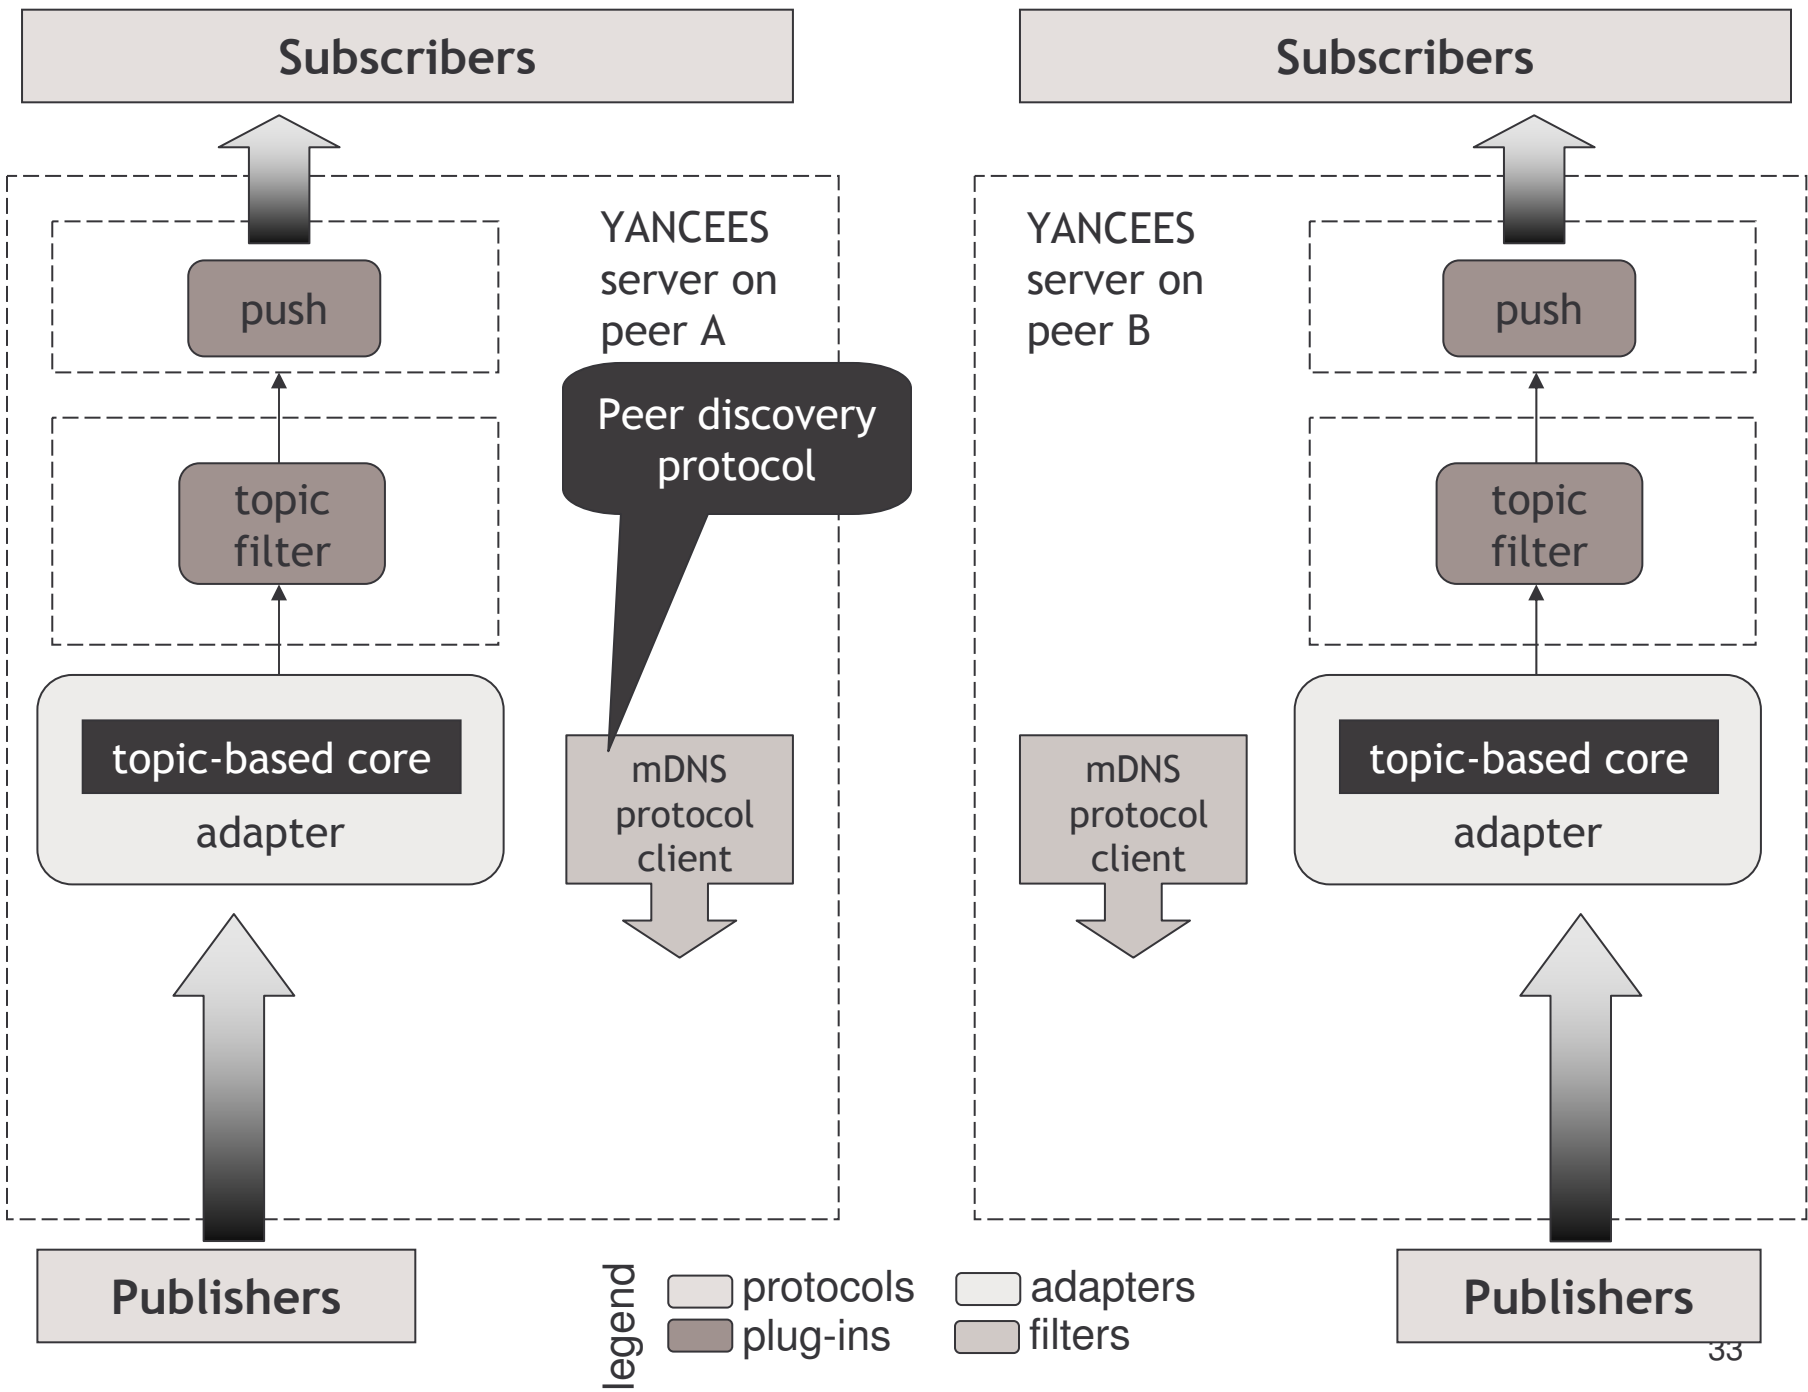
Publishers

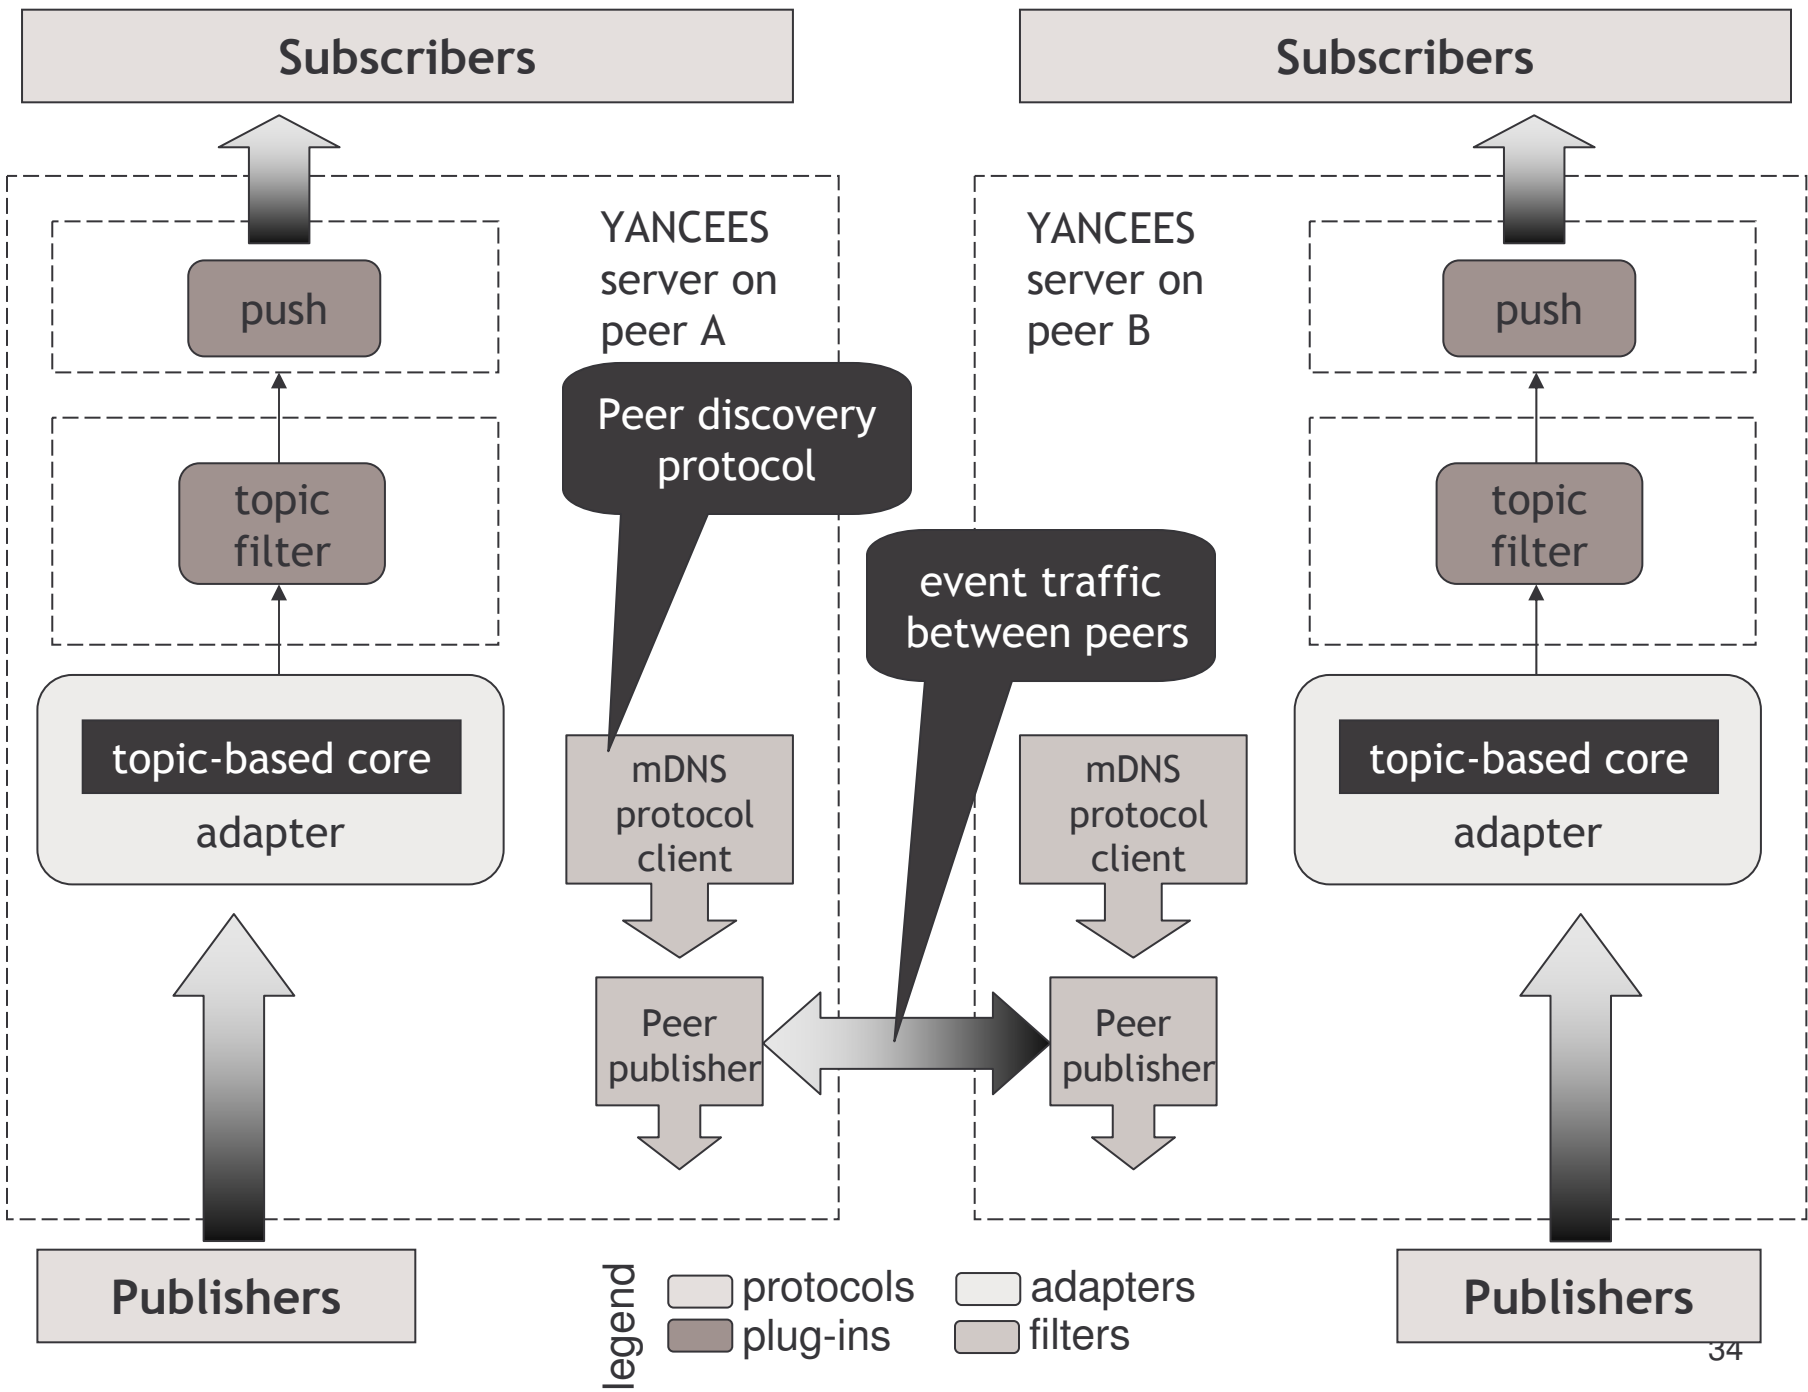


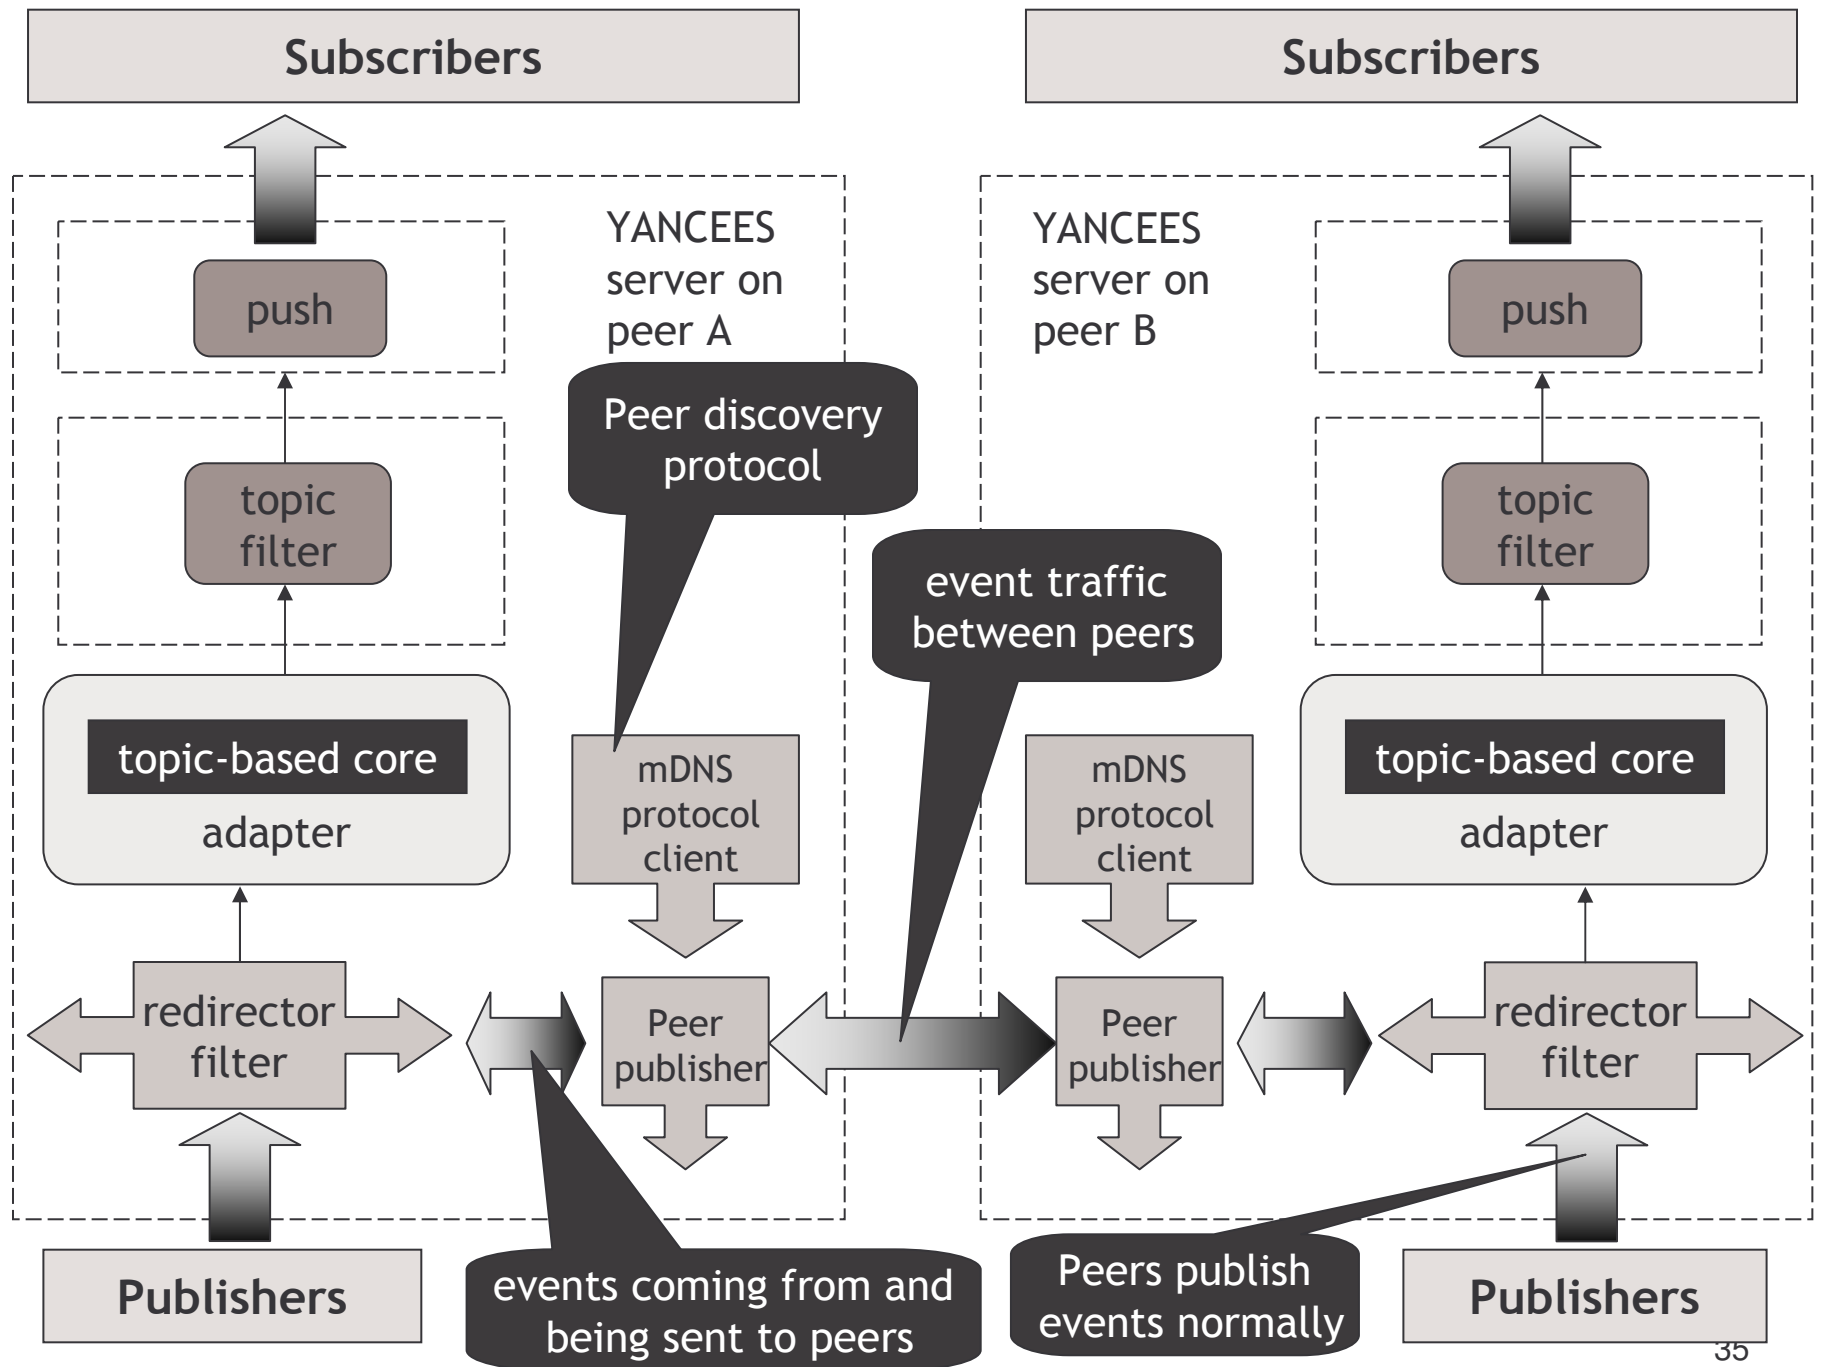


legend

	protocols		adapters
	plug-ins		filters







Software and security visualization

- YANCEES addresses two different routing requirements:
 - Software visualization: support fast routing
 - Security visualization: content-based filtering

Software visualization tool

The image shows a software visualization tool monitoring a web browser. The tool's interface includes a 'Method depth' window and a 'Color Chooser' window. The 'Method depth' window displays a visualization of the program stack, with a callout 'Loaded classes' pointing to the top part of the stack. The 'Color Chooser' window displays a list of classes and objects, with a callout 'Object hierarchy in memory' pointing to the list. A callout 'Program stack visualization' points to the main visualization area. The background shows a web browser window displaying the DOJ homepage.

Clue v4.2: DOJ: U.S. Department of Justice Home Page

File Options Bookmarks

Address <http://www.usdoj.gov/>

Clue WBC 4.2 evaluation copy -- for personal use only. Copyright (c)2001-02 NetClue Corp. All rights reserved.

Accessibility Information

United States De
JUST

HOME PAGE | CONTACT US | P

What's New

New On Site
Components' News
What's Hot

About DOJ

Alphabetical List of Components
Organization Chart
Budget Information
General Information

Publications & Documents

AG Annual Reports
Performance Reports and Plans
Legal Documents
Reports & Pubs: Alpha - Comp Lists
Strategic Plans
Information Quality Guidelines
U.S. Attorneys' Manual

Employment

No Fear Act Employment Data

FOIA

Status: 01whatsnew/01_3.html

Method depth

Loaded classes

Object hierarchy in memory

Color Chooser

- ClassAnalyzer
- Event
- ClassProperties
- MenuShortcut
- ShortcutHandler
- ComponentEvt
- NativeGraphics
- Graphics
- FocusEvt
- WindowEvt
- ActionEvt
- security
- net
- kaffe
- generalTest
 - main
 - <init>
- SymAction
 - recButton_actionPerf
- RecThread
 - <init>
 - tailRec
 - headRec

Detailed Sparse

Program stack visualization

Application being monitored: web browser

Security visualization tool

The screenshot shows a web browser window titled "Clue v4.2: DOJ: U.S. Department of Justice Home Page". The address bar shows "http://www.usdoj.gov/". The page content includes the DOJ logo, navigation links (HOME PAGE, CONTACT US, PRIVACY POLICY, SITE MAP, SEARCH), and a "What's New" section. A "Network Activity Visualization" window is overlaid on the page, displaying a list of connections with their status, throughput, and bytes sent/received.

Connection	Activity	Throughput (KB/s)	Bytes Sent	Received Bytes
localhost:3272-www.usdoj.gov:80	Closed	4	45 bytes sent	11,683 bytes received
localhost:3273-www.usdoj.gov:80	Closed	0	130 bytes sent	230 bytes received
localhost:3280-www.usdoj.gov:80	Closed	0	205 bytes sent	229 bytes received
localhost:3279-www.usdoj.gov:80	Closed	0	236 bytes sent	224 bytes received
localhost:3281-www.dhs.gov:80	Closed	0	117 bytes sent	412 bytes received
localhost:3282-www.usdoj.gov:80	Closed	0	88 bytes sent	339 bytes received
localhost:3283-www.usdoj.gov:80	Closed	1	240 bytes sent	14017 bytes received

History of active/inactive connections, and their information flow

Java web-browser dynamically instrumented by our infrastructure

Java Runtime Environment

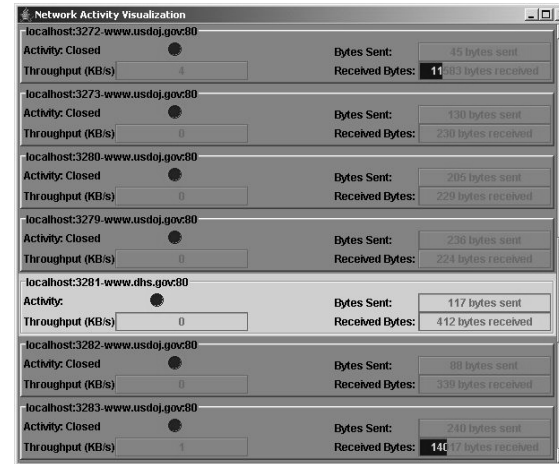


dynamically instrumented application

Vavoom Class Loader

Application .class files

network activity visualization




subscription

filtered events

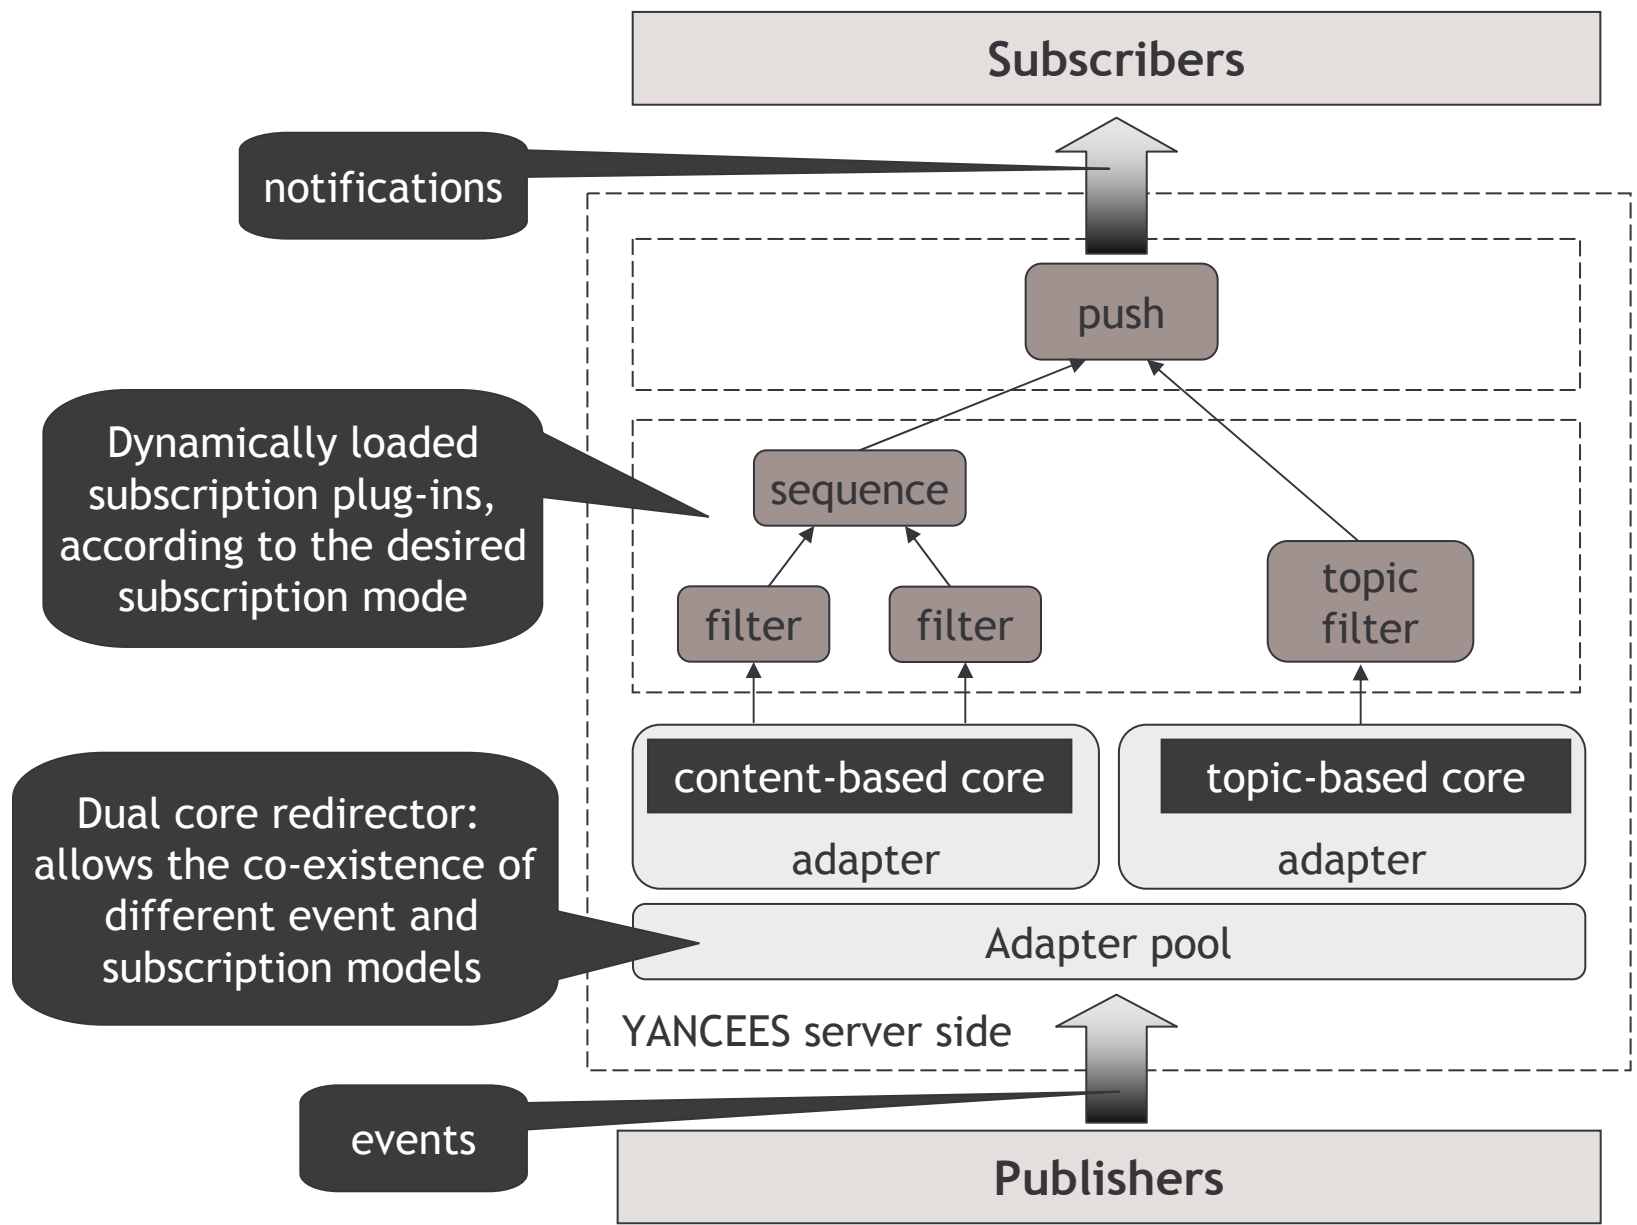
YANCEES
Publish/subscribe event router

publish execution events

routed events



Software visualization



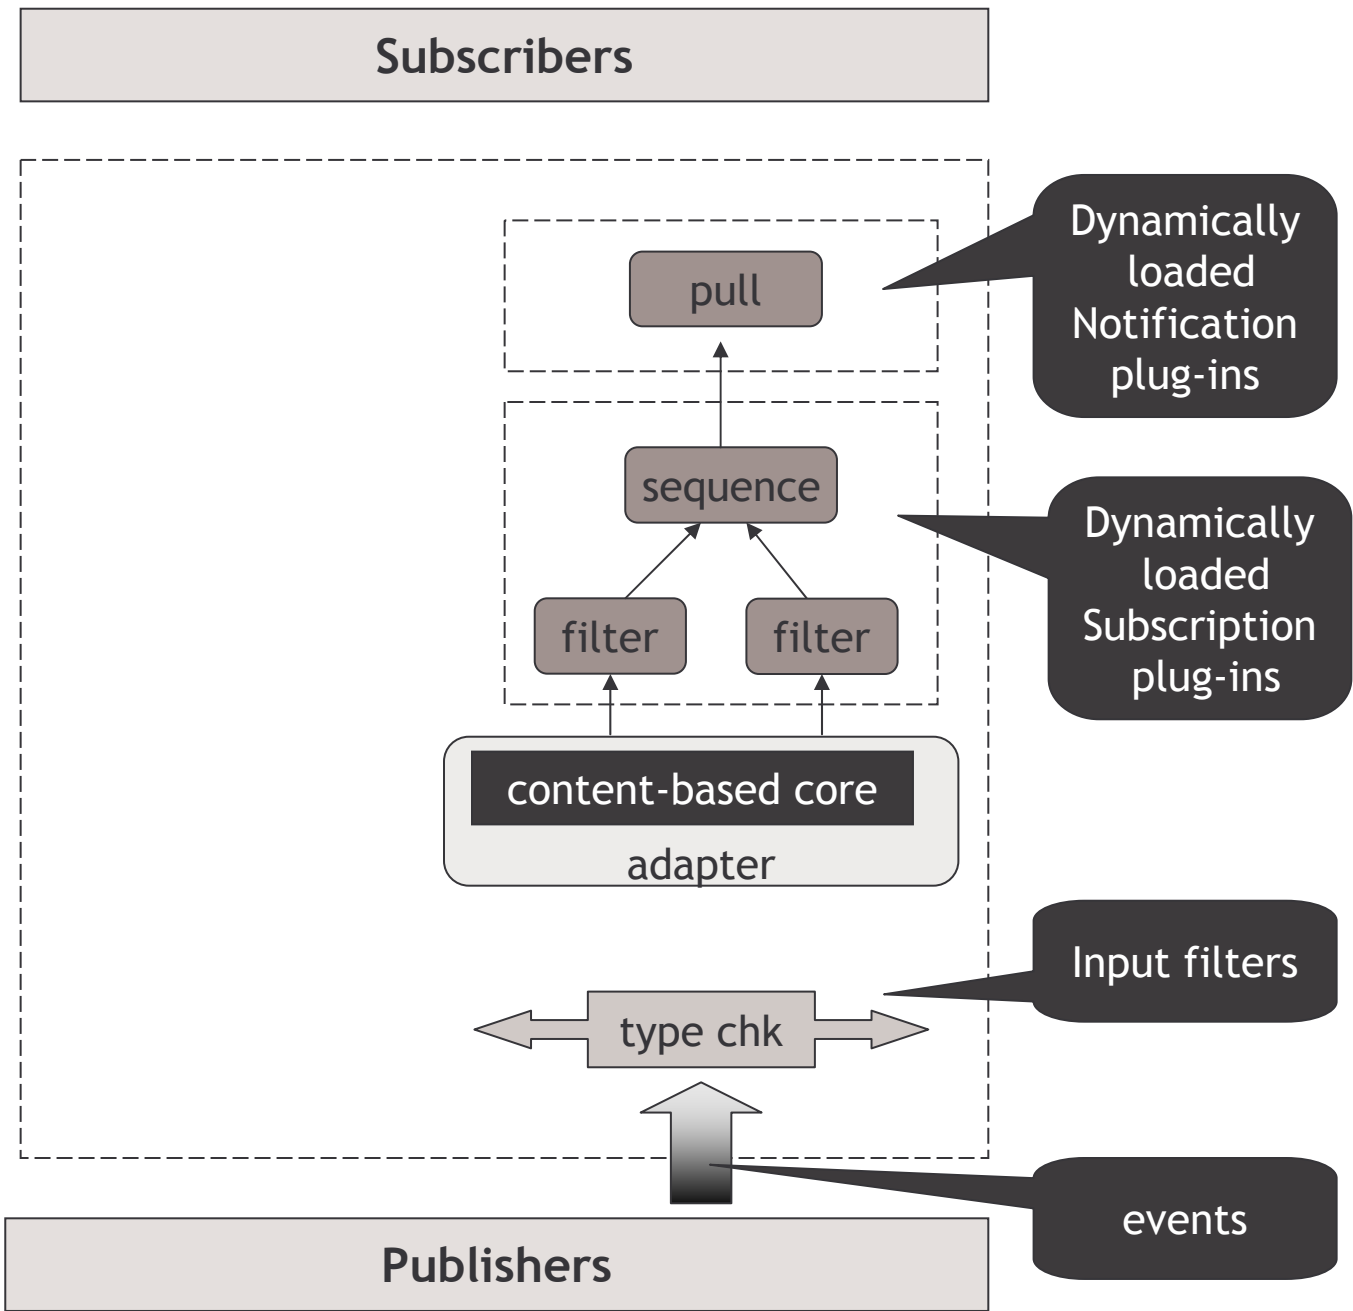
Dynamically loaded subscription plug-ins, according to the desired subscription mode

Dual core redirector: allows the co-existence of different event and subscription models

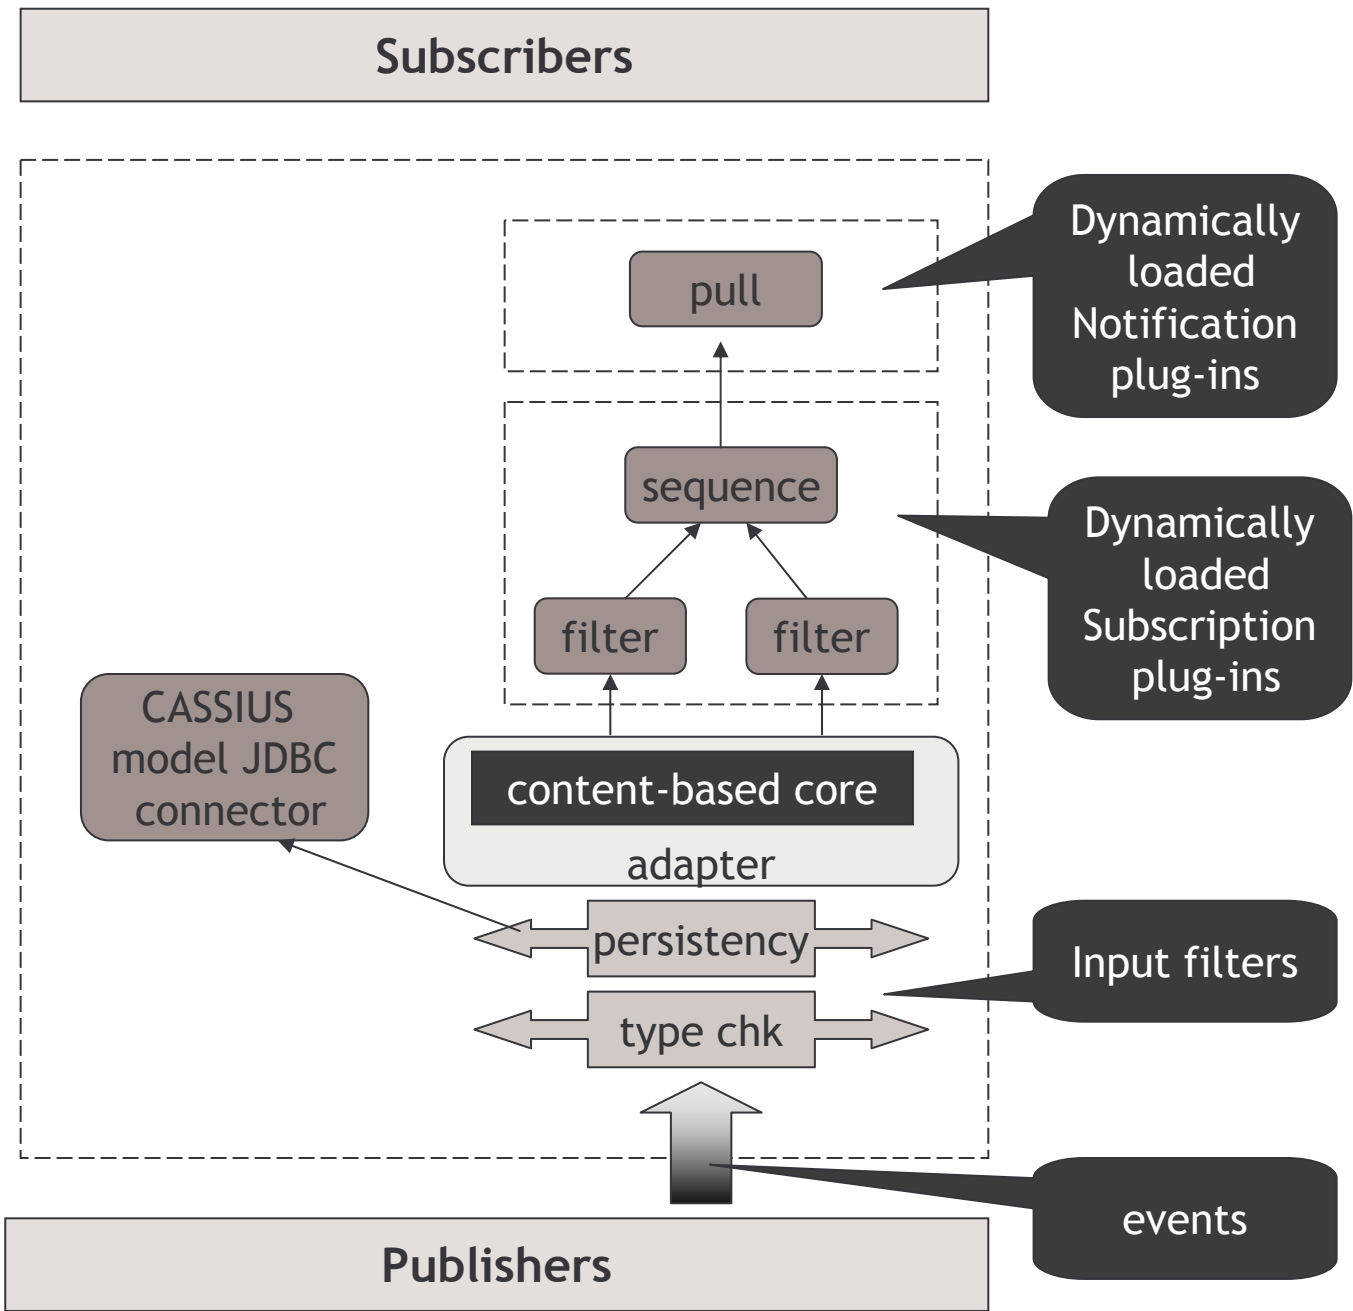
legend
 light gray box: protocols
 dark gray box: plug-ins
 light gray box: adapters
 dark gray box: filters

CASSIUS – awareness publish/subscribe server

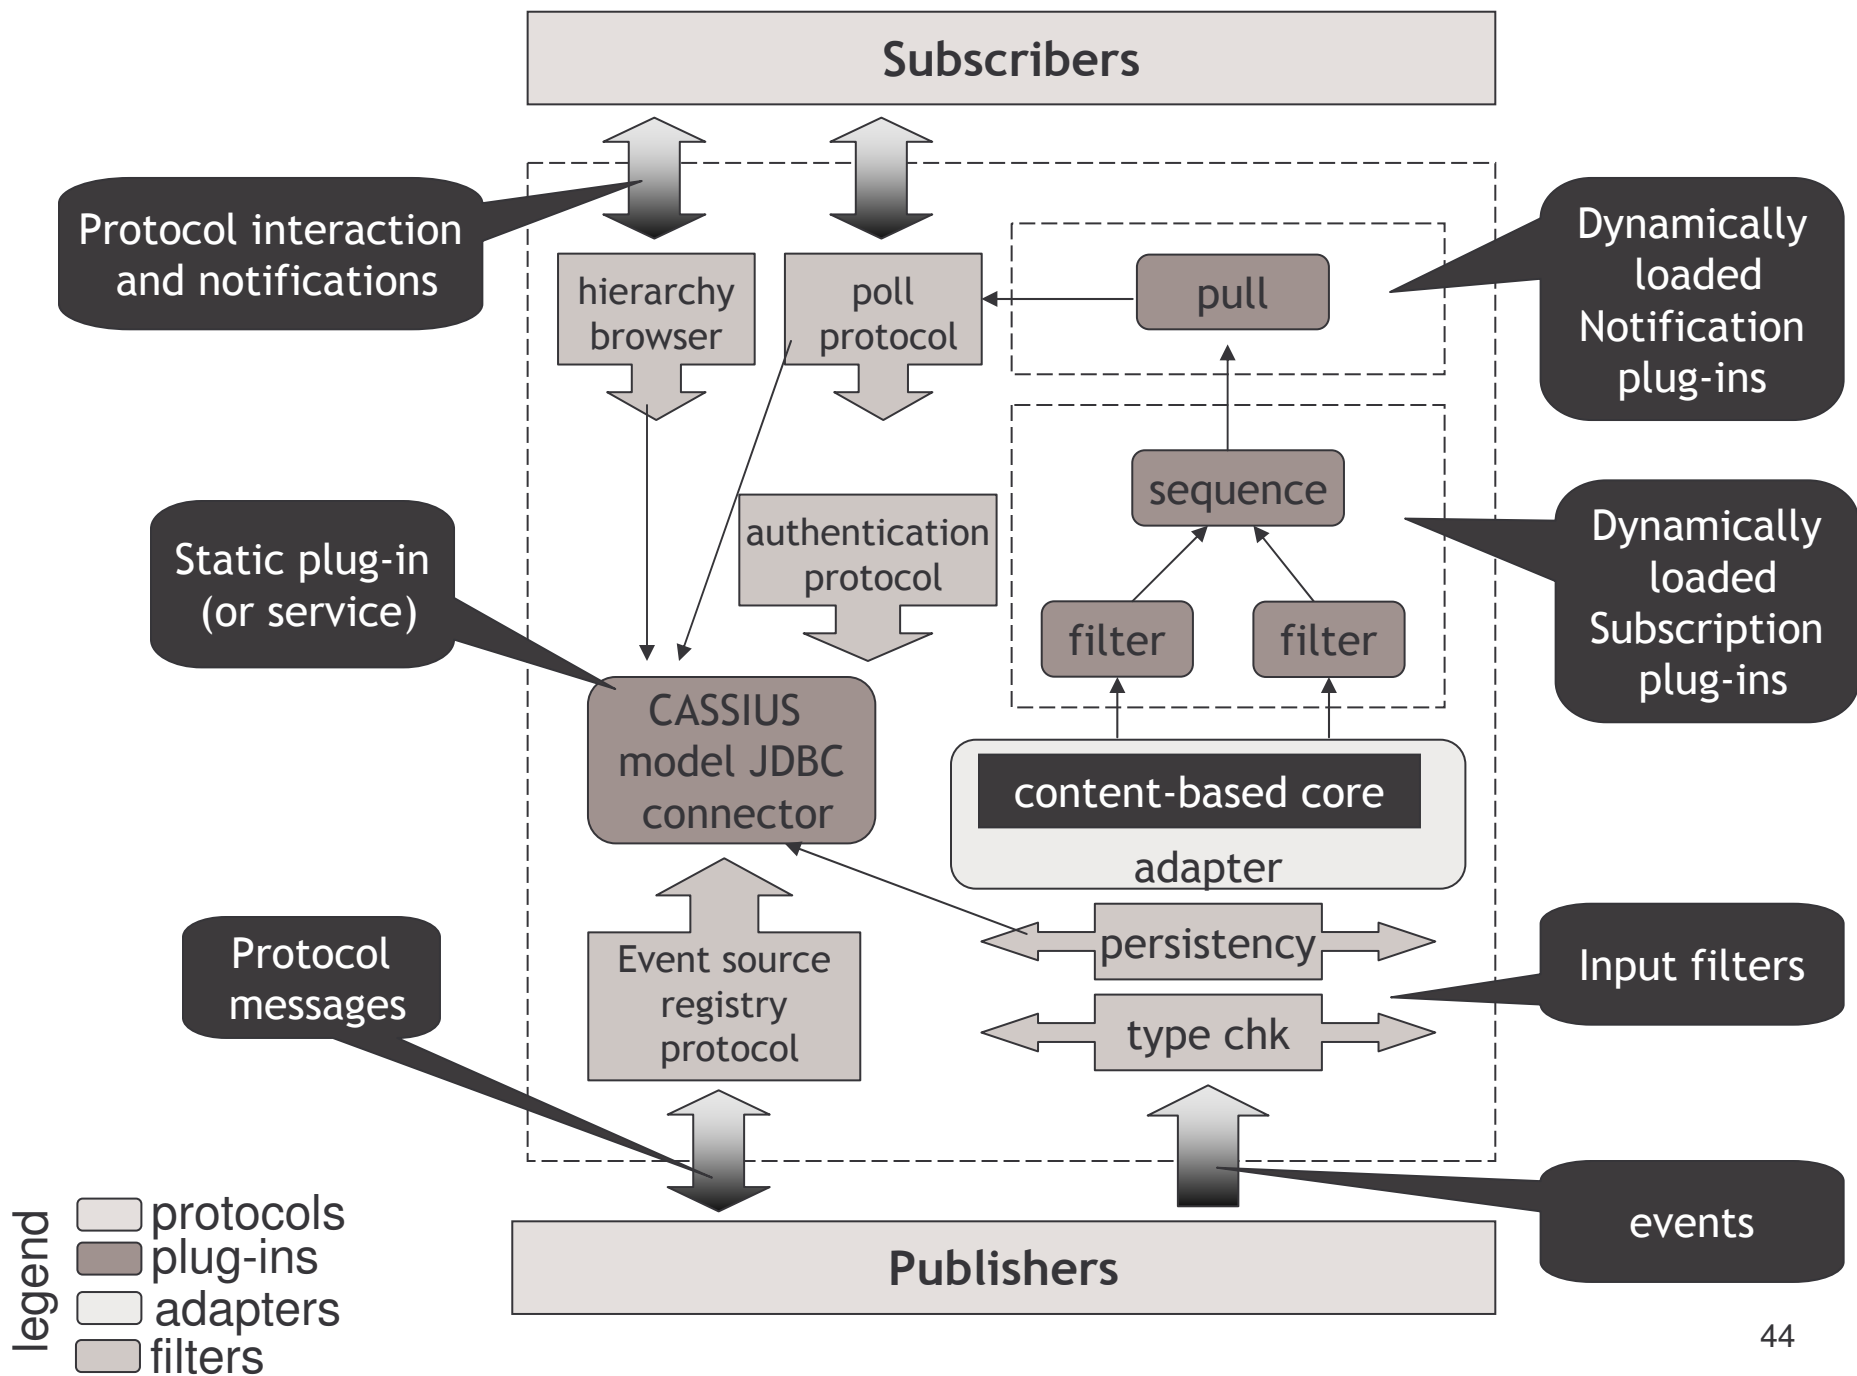
- CASSIUS pub/sub model provides:
 - Event persistency
 - Event typing enforcing
 - Pull notification delivery
- CASSIUS also supports the following protocols:
 - Event source discovery
 - Event type hierarchy browsing
 - Authentication



- legend
- protocols
 - plug-ins
 - adapters
 - filters



- legend
- protocols
 - plug-ins
 - adapters
 - filters



Conclusions

Advantages of the approach

- **Configurability:** The combination of **plug-ins and extensible languages** provide coherent composition of interdependent features;
 - the subset of language extensions and plug-ins also define the **footprint** of the server.
- **Extensibility:** new features can be provided by extending the language and implementing new plug-ins and filters
- **Reuse:** plug-ins can depend on one another, speeding up the development process
- **Support for multiple infrastructures:** the microkernel approach allows different publish/subscribe cores to be installed at the same time
- **Variability:** plug-ins can be installed at load time (configuration file) and runtime (downloaded as needed). They are also allocated according to the application needs
- **Multiple event models:** adapters to different pub/sub cores permit multiple event representations to co-exist.

Drawbacks

- Performance:
 - In our experiments, the XML technology (subscription and notification parsing) adds an extra 100 ms to the subscription process (but this is a one time cost)
 - The plug-in hierarchy adds an extra 50 ms to the notifications routing time (but the throughput is compatible with Siena and Elvin ~8000 events/second) due to our buffering strategy

- Framework costs:
 - Initial generalization and implementation
 - Initial learning curve (not much worse than more advanced pub/sub systems as CORBA-NS)

- Non-functional requirements are not so easy to implement (need to extend many points in the system, AOP may help)

Future work

- Address usability issues
 - Achieve a balance between model complexity and its extensibility
- Study the use AOP for non-functional requirements
- Study the use of rule-based patterns for more complex event processing
- Perform usability case studies

Questions/Comments?