

1. Meeting Information

Date/Time of the Meeting:	Mon Jan 23rd 2012, 10h (Delayed as some partners not yet here)
Inviting person:	
Minute taker:	
Name of the meeting:	Data/Context Assets Status
Place of the meeting:	East 1, floor 0, room C61
Phone details (if PhC):	
Version	

2. Attendees

Name	Company / Organization
Chema Fuentes (Tue-Thu)	Atos
Remi Challamel (Wed & Thu 15h)	Thales
Carlo (Tue 10:30 onwards)	Telecom Italia
Peter (Mon 14h onwards - Wed) Andreas (Wed & Thu)	Siemens
Fano Ramparany	FT-Orange Group
Tali Yatzkar-Haham	IBM-IL
Ken, Grant, Gregorio, Carlos	Telefónica I+D

3. Minute Taker (optional)

Meeting	Minute Taker	Date
ALL	Carlos Ralli	Telefónica I+D

4. Objective of the meeting

Data/Context WP meetings as schedule in the overall agenda.

5. Topics discussed and Agreements

Id	Topic
1	Monday 11:00-12:30 plus 14:00-14:30) CEP - BigData - PublishSubscribe integration SAMSON-AMIT integration (presentation) - Scenario1: where AMIT processes incoming “buy” events and Queries MongoDB Samson Database to ask if someone is “a Barcelona team fan”. As a result, AMIT will obtain as output events a campaign targeted customers. Motivation of the scenario: business logic may change over time (i.e change Barcelona for R.Madrid or select certain fans, etc). Roadmap to make things work in such scenario1: A) Customer side: Define the specific scenario and provide it to Samson/AMIT

teams.

* For Samson: "Customer" to define the scenario in terms of the format and example of the data/profiles to be collected. Also the live data source should be established.

(Samson Incoming data collection: read from a Database, TCP/IP socket, file, library for pushing data to Samson.)

* For AMIT: "Customer" to provide business rules (examples like Scenario1 definition) and events filling up some forms. There is an event processing language. Database query is one of the expressions in the language.

Tali performs a Demo of Amit UI:

- Event form to declare its attributes, type, etc
- Definition of new rules.

B) Samson team side: start getting data (format and samples of the data to be collected. It could be: CDRs, WEB logs, etc), programming the analytical module (Samson),

C) AMIT team side: provide support to the customer so he fills the forms up, connect to the live data source (Incoming events) and connect AMIT output to the customer facilities

How we could advance on integration so the "customer" is exposed to one single interface?

Amit would be exposed to the customer and Samson would run in the backoffice. Customer accesses still the AMIT forms but a middleware might establish AMIT-Samson interconnection (database list within AMIT).

Amit needs to know the database name, the IP where it is running and also the structure/name of the fields (this should be known by the middleware handling the query and the customer defining it).

Next steps: Amit to be able to connect to MongoDB. Tali: seems to be feasible. Samson would need to know basic data types to manage.

Scenario (II): Samson output used as one more source for Amit incoming events.

APs:

- Think on a specific business use-case (example of applications: whenever somebody -fan of a club- enters a specific cell - stadium- we trigger an action/event to be provided to Amit).

Incoming events: specific adaptors and generic adaptors (they will be specified in D2.3)

- How to filter events that Samson does not overload Amit => Scenario II++
Analytical module, more complex but not to be built from scratch everytime we change the applications.

	Scenario (III): Amit feeding Samson. Perhaps to change the campaign of Scenario I.
2	(Monday 14:30-17:30) Baseline Assets Status: 2.1 Samson Presentation (TID) - Near RT analytics with Streaming MapReduce. - Multiple input sources, static or continuous - Output state (results) can be stored into static files of managed storage (MongoDB or HDFS). - MongoDB: medium sized data sets 100-500Gb - Hadoop HDFS + Hive: SQL interface, large data sets > 0.5 Tb Sprint review: - S1: integration discussions with CEP GE. - S2: packaging . Done, needs some testing. - S3: Devel of virtual machines (more or less done), Architecture spec (D2.3) - Next steps: redundancy, de-centralized control, fault tolerance, improvements in streaming. - Requirements: CEP integration discussions, middleware/REST discussions. Sprint 4: Refine the roadmap that have been influenced by FI-WARE discussions. 2.2 CEP Presentation (IBM) - From event-condition-action to pattern-condition-action - In some scenarios, single events are insignificant but situations (complex combinations of events) are. - Rules are established before the runtime, where incoming events are processed and detected situations are triggered according to the definitions (rules). ACTIVITIES: A1) Proton Adapter framework (evolution for Amit, but many things are built from scratch) Adapter Framework design (u-s, done), adapter framework skeleton implementation (u-s, done), file adapter implementation (u-s, nearly finished). - Next steps: Architecture doc that describes this adaptor framework (u-s), Rest

adapter (u-s).

A2) Also Proton - Build Time Web based UI

- Web UI fm (u-s, done)
- Next steps: Definiton Forms Implementation (several u-s, implementation started).
- (Events, Context, App Network, Agents -> validation -> JSON object = input for configuration).

A3) Integration Discussions

- WS with BigData (w-i, done)
- Initial discussion with IoT (w-i, done)
- To achieve in Mad meeting: explore integration with BigData & PubSub.

2.3 PubSub Presentation - Semantic engine (Orange)

- RDF store mirrors Context data Cache to make the Semantic engine able to answer SPARQL queries. RDF store is more flexible whenever you want to add and access new classes/attributes (this can be done at runtime in the RDF store, but not in the Context Data Cache).
- An ontology mapping defines the semantics of the CML (Context Modeling Language) (i.e notify me when room 24 is lighten up)

Components details:

- Ontology mapping (to be developed)
- Semantic Ifc Mapping (to be developed)
- Semantic engine (commercial products available. e.g. virtuoso)
- Context switch (to be developed).

Semantic Interface functionality: enables RDF modeling of data/contetx (done), supports queries and subs expressed in SPARQL(done), provides RESTful access point (under development).

Unstructured data to Ontology compliant RDF translator.

Sprints 1,2,3 in line with expectations.

Main Next Step: Integration with TI asset scheduled for Sprint 4 (end of Feb).

2.4 Multimedia Analysis (Siemens)

- Enabler for MM analysis and information extraction
- Automated filtering of relevant or critical events in Media streams (person detection/tracking)
- Baseline asset for compressed domain media analysis: codoan

	Example of change detection: indoor video surv, inspection of RTP packets only, regular PC can scan ten k of streams. Normally you don't need to go into the stream but just the packet size (for lights on/off, person detected, etc.) - Moving object detection: not necessary to decode the full stream. Resolutions does not impact so much (linear increase of complexity with images resolution). MM analysis Backlog: -S1.1.1: Meta-data format (w-i, done) -S1.1.2: event detection notification (u-s, under execution), Arch. Specification (w-i, under execution) -S1.1.3: same as 1.1.2 (continuation). 2.5 Metadata Preprocessing (Siemens) - Encapsulation of transport and metadata transformation as-a-service usable from other web apps or components. - Generic metadata transformation approach. - Metadata stream processing. Main development for this asset was not plan for 1st minor release (Sprints 1, 2 & 3).
3	Apps - Data WPs meeting (Tue 9-10h) Torsten-SAP, Tali-IBM, Chema-Atos, Axel Fasse-SAP, Gregorio-TID, Ken-TID Target of the meeting: how to integrate Data/Context in the business framework. Apps team: integration of Data asset should be quite similar to the guidelines to be provided to any other kind of component. Torsten Presentation - Overall Architecture: Aggregator, Channel Maker, Broker, Gateway. Roles: Consumer, Provider, Premise and Host. - Aggregator (Composition and Mash-up). - To include components in the Broker (Business framework), Describe the service in a common format (business aspects: i.e pricing model, etc.). - Channel maker (channels & devices): Services can be delivered to multiple channels.

	- A Gateway (interoperability) performs some mediation as long as components will use different technologies. High-level architecture/overview of the Business Framework: - Revenue sharing, SLA management, USDL repository&registry, marketplace and store. - BE & BM provisioning is open now and it is actually part of the 1st Open Call. Things to take into account for the integration: (->Chema ATOS input here) - Components have to be described in USDL. - Aggregator has 3-4 ways of composition. There are 3 mash-up platforms, to analyze which one will be used . Let's work jointly analyzing how one data component could be integrated into the business framework. - Gregorio provides a description of BigData asset. It is actually more targeted now for developers. - 1st think on the kind of interface/functions we want to expose from this specific component. Not a problem it is oriented to developers, focus on what it is offered. - There are different services in this Bigdata description: one simple example might be to provide a data stream and get as a result the analytics. One specific example of service that might be offered: an operator provides a live stream, list of users and a set of locations and gets in turn events triggered whenever those users enter/approach those locations. - Things to work out in this example: What are the interfaces and parameters (REST), pricing model (possibly described in USDL too) and USDL descriptions of components. Conclusion of the exercise: DATA Components interested on this Business framework integration should first think on the specific services to be offered and the Business model (steps above). The idea is to work in parallel to have the detailed description discussion once Apps chapter gets the framework ready too.
4	Baseline Assets Status (Cont). Tue 10:00 (+30 min delay) 2.6 THESEUS Query Broker (Siemens AG) Core modules are implemented - Comm with QB currently by valid MPQF queries - Connection of data repositories

Example Scenarios:

- Interoperable image search (image similarities), medical decision support for radiologists (text/parameters associated to images)

Roadmap: basic core package, interface via API calls, support for REST/SOAP communication, intra-query optimisation ... (see slide)

Sprints: No major deviations, minor impact of Xmas.

1.1.1 flickrAdapter (u-s, finished) Retrieve images based on keywords

1.1.2 Core Package (u-s, under exec)

1.1.3 Core Package (u-s, under exec), Arch spec (w-i, under exec)

Question arises: JDBC interfaces to be supported? Not really, why is it relevant? For CEP-IBM it might be interesting. Then we can discuss this feature further.

2.7 (Telecom Italia) (?)

Context Management Services PublishSubscribe (Telecom Italia)

Context Query and Context Query Language (CQL)

Motivation: Context-aware apps need to query the context (raw context info, asynchronous notifications, advanced filtering and conditions on multiple context info).

Overall architecture: "Local context broker" collects context from user devices (Cell ID, device info, movement, speed, bluetooth ...), "context broker" performs the aggregation processing in RT this contextual info. Restful API with the Applications&Services enabling to retrieve context info from them and also whenever any contextual situation is met.

Contextual model: context info is represented in "ContextElements" describing a certain "ContextScope" (group of "ContextValues") for an "Entity", with associated "MetaData".

ContextValues and Metadata can have multiple representations.

An example of XML binding is shown.

Context Query Language (CQL): XML based, main tags: "entities" with entity types, "scopes", action to be performed "select"/"subscribe" and conditions with condition times "ONCLOCK" ...

Using CQP apps can specify filters and constraints on both context data and metadata.

	Notification: Implemented methods: HTTP POST, Binary SMS, XMPP, Apple Push Notification service. No deviations in the current sprints.
5	Testbed Requirements 5.1 Context Aware Platform for Pub Sub - Runs on JBOSS platform There are instances running on Linux/Ubuntu or Windows 2003 and 2008. CentOS Number of Machines/VM: a single machine would be fine, double core, 4-8 GB RAM. 5.2 Bigdata Analysis 64-bit - Linux/Ubuntu 11.04, Redhat 6, CentOS 6 (Ubuntu preferred) One machine would be enough (proof of concept), 4 or more cores. 2 Tb Disk, 4-8 GB RAM. For performance 3 or more machines are required For the scenario1 (CEP+BigData) another machine for MongoDB server might be useful. Additionally another one for Hadoop. 5.3 Semantics for PubSub For Vituso if would be wise to have another machine with 8-16 GB RAM 5.4 Multimedia Analysis 1 computer dual core, Windows XP onwards - 4 GB RAM 5.5 Metada pre-processing Similar to 5.4 except .NET 4.0 framework is needed in addition 5.6 Query Broker Similar to 5.4 too. Java JRE 1.6 is needed in addition 5.7 CEP Version 1: JVM, Tomcat server, Webserver for REST adaptors, Java messaging infrastructure (if JMS adapters are used) 5.8 Semantic App support GE 1 machine JBOSS, Linux, Sesame repository (8-16 GB RAM, dual core, ...)

	SQL database (ontologies storage) to be run in the same machine.
6	Architecture Discussion ----> Rest APIs Roundtable - Semantic Apps Support (ATOS). Rest APIs will be implemented when its useful (in the roadmap and under execution). Easy to implement using Java (Jersey reference). Similar to deploy a webserver in Tomcat. Not intuitive for some specific functionalities. -> Multimedia analysis and metadata preprocessing should be easier but more difficult for the Query Broker. - CEP: a Rest adaptor is planned (already started to work on it). In the standalone service there are no webserver so it is less natural but can be done. In the J2EE it should be more natural. For the rules currently we use the WEB forms UI but we might study an API for providing the JSON expected file (however we do not see it necessary and will not do it unless there is a specific case). To ensure assets interaction will work we should agree first on JSON or XML. If XML on on the schema and tags of the XML files to be exchanged over the Rest APIs. We might follow the example of SPARQL protocol RDF. XML parsing is slower than JSON but from the size point of view they are almost the same. Schema and potential functionalities are better in XML. JSON schema possibilities is weaker. - Functions common to all assets: 1. Status ("0 - OK" , Error String -including code-). 2. Reporting programmed tasks (Init, halt/shutdown, overflow, supervision mode ..) 3. Basic Control Interface (Start, restart, ...) ** BigData: - Monitoring (1, 2) + list of buffers + list of operations (methods) + memory usage/CPU usage - Control (3): start/stop an specific script - Data Stream Channel handshake (customer asks possible methods and Samson answer: this file or this IP:port or ...) - Event Flow Registration: The customer applies for a specific type of events triggered by Samson resulting from the data stream analysis. E.g. "everytime a user enters home". To interact with the CEP: xml file with event name and event attributes.

**** CEP:**

- Monitoring/Control should be feasible for later versions
- Get events.
- Send events.
- Deploy the Rules definition file/version. Input JSON file/version.
- Get the Rules definition file/version.

**** PubSub:**

- Monitoring/Control should be feasible. Suggestion: use standard HTTP codes.
Monitoring (specific): Get list of Context Providers, list of support entities, list of supported scopes
- Get/Update/subscribe context
- Subscribe/Notify (notifications)

**** Multimedia Analysis:**

- Monitoring should be feasible.
 - Control: Start/Stop
 - De-/Register the source
 - De-/Register the consumer
 - Configuration of processing (which events should be processed)
- Media stream would be RTP.

**** Semantic :**

- Get list of deployed ontologies (name+version)
- Get/Post/Delete/Push Ontologies into the registry

Quite nice functionality for a middleware would be to manage the temporary messages for a period an asset is not available.

----> Middleware:

As long as most assets are planning REST APIs no middleware is needed for interfacing with 3rd parties.

Perhaps a middleware would provide a "Distributed database"?

---> "Dummy use-cases".

- Samson can process navigation logs of the users (userID+log) (millions of users or an ISP) and derive conclusions like what's popular now? or what's popular for a specific user? etc.

- Query Broker allows to find out medical images (X-ray) related to a disease or to specific people/ages.

- Multimedia analysis allows you to determine what people/objects are coming in/out to a room in a video file/stream.

	Combination with CEP: people appear in certain rooms during the night -> check if a guard visited some rooms N times. Another case: check if there are more people than the guard himself! MM Analysis would be a context producer (this room is not empty, a guard is in this room, etc). CEP would be context consumer, so only receiving events when there are changes in the context. Final dummy-scenario: surveillance cameras within a stadium area. Control/monitoring of guards/people as described above plus Bigdata analysis: retrieve the list of an ISP users nearby that stadium in the event of emergency. Ontologies and workspaces regular updating.
7	Tue 16:30-17:30 CEP and PubSub integration discussion - Discussion framework: Events associated with entities, vice-versa is context. We do need context, we do need entities and we do need events Event = Event name + list of attributes + timestamp When sending the event we might associate it with an entity or not. However to be processed by Context Management it has to be associated with an entity. CEP might add a new attribute "entity" or maybe a special adaptor could be done for PubSub interface. TI component expects: Entity type, entity ID, scope of the context, value of the attribute NGSI topic (TI): being compliant to NGSI means no great thing because there is no actual reference implementation and it relies on very abstract information models. The most of NGSI is optional. Compliance is quite hard to be defined. We could try to agree internally on a kind of interface to be exposed to IoT and compare with the one worked out by IoT (it seems they have some previous work going on).
8	

	Some notes during IoT-Data meeting (these are not this meeting minutes!) 1. IoT Data model (Telefonica) 2. NEC presentation - OMA NGSI model Context Information Model. Entity <-> Context attributes <-> Metadata Entities are: Person, Group, Place, Event, Thing (Juanjo: Data/Context WP proposed abstract model is compliant with the definition above - see FI-WARE vision in the Wiki for more details) NGSI-10: Operation APIs (Update, Query, Subscription) - Update: update a context value in the Context Management Component (CMC) - Query: sync/one-time access - Subscribe: async/continuous access NGSI-9: Federation APIs (Register, Discover/Subscribe) Context Management Systems need to federate (i.e a building may provide info to the operator. Note: NGSI-10 can work in both directions so operator can ask building) 3. TI-presentation: Context Management Services NGSI-10 is supported but not all operations of NGSI-9. To be checked with the rest of TI team. TI to send some additional input on the NGSI discussion aprox. 1 week after Madrid meeting. RESTful binding will be sent again by IoT WP Leader. 4. CEP presentation (IBM-IL). Overall presentation of AMIT/Proton asset (IBM Event-based Middleware & Solutions group).
7	Wed 14:00, UC Project tickets backlog CEP, PubSub, BigData, MM + metada preprocessing + Query Broker, Semantic App support, Location Server. - CEP: tickets are usually Data specific but not CEP specific.

	Two tickets to be assigned to Massive “Data gathering” (done) or “distributed database” (pending). 4 tickets related to CEP but not just to CEP. Thales: a 3r party may activate a tracking process (location every x minutes). Ticket. 321 to (Location Server + Context management) Ticket.393. Related to I2ND as well (perhaps we might add geo-location if needed or store all this as contextual information). Go back to the UC project asking for clarification (example(s) and story-line(s)). (Make an explanatory session for NGSi within Data).
8	Thu 14:00 Assets status review 2.9 Location Server Platform (Thales Alenia Space - TAS) - Manage mobile end-users location retrieval Most recent protocols: MLP3.3 (iface with location systems), SUPLVv2 (iface with IP hantsets). Mechanisms: Cell-id, Wifi, A-GPS - Key developments so far (Sprints 1,2 3): None - Deviations: Developments delayed for April 2012 (Sprint 1.2.2) - Review of tickets 313, 321, 457 and 393 (from UC projects). 393 to Carlos 321 Thales and TI to organize a Workshop (en Feb). Minimum outcomes: interfaces and dummy-use cases (covering at least UC project ticket 321). 2.10 Semantic Application Support (ATOS) RDF/OWL storage and retrieval capabilities. Workshop to be done for integration with MM Analysis+Query Broker+metadata pre-processing and semantic contribution (Orange) to PubSub Asset (TI). ATOS to clarify what service/added value they may provide to those assets above and conclude together with them on the interfaces and dummy-use-cases. S1: Decisions adopted (work-items). - OMV (Ontology Metadata Vocabulary) selected as ontology metadata format. - Ontology registry definition. SPRQL - Ontology storage evaluation. Ontologies will be stored using JPA. S2 and S3: defining and building Rest interface: Ontologies.

	UC projects tickets: interesting 6 tickets from Safecity (2-Envirofi and 2-smartagrifood). All of them are more or less relevant for the asset even they are too general or too domain oriented. Architecture board feedback: many UC projects interested on streaming analysis (perhaps BigData + MM Analysis). Dedicated mailing list for this: "fiware-streaming".
	Thu 15:30 D2.3 - Architecture Document deadline: Feb 17th 2012. We should understand this document as "Architecture of the interfaces" rather than "internal components architecture". It is oriented to programmers as main target audience so they will be able later on to connect applications to our components and invoke their operations. WPI agreed not to provide any binding such as Rest bindings M12 Deliverable will actually be a continuation/extension of D2.3 "Programmers guide" deliverable will be mainly tutorials. Use "FNC block Diagrams" notation for the architecture images (buiding blocks), not for other kind of images. Check the e-mail from Juanjo forwarded right today by Grant (with links to tools but you can manually create too). Query Broker is in a quite good state and might be used as a reference for other assets. However, Juanjo is to send an example for PubSub. - Query Broker is in a quite good shape according to D2.3 expectations. In all what respect to Interfaces the arch specs should elaborate which are the interfaces, entities and names of operations (no more details are required beyond the name). - Metadata pre-processing (now refactoring the component so documentation is on its way. It should be ready 1-2 weeks for the names of operations). - MM Analysis: 80% in the Wiki but needs some more work. - PubSub: not yet started, They will check also the example on PubSub to be provided by Juanjo and see if there is a way to complete it further. - BigData: It's on the way and aligned to expectations. It will be finished by next week.

	- Semantic Application Support: Not yet started. - CEP: There is a document (shown by Tali in the meeting) describing Proton asset (formerly known as Amit) adaptors . Needs to be uploaded to the Wiki (it should be done in 1 week approx).
	Thu 16:30 Next Steps: - Analysis of UC projects tickets - Next Workshops on Data assets integration: LS-PubSub, SAS-PubSub+MMA+QB+MPReP

6. Summary of action points

Id	Action point	Responsible
AP-1	CEP-BigData integration: Scenario1: Amit to be able to perform MongoDB questions. Also to provide Samson team with the basic data types. It seems feasible to plan this development for Amit (forthcoming Sprints). IBM to come back with roadmap and possibly doubts/questions.	Tali, IBM
AP-2	Scenario II: - Find out business use cases together with IBM team. - REST APIs: CEP for Feb'12 and SAMSON to provide a roadmap for REST APIs	Telefonica

AP-3	AP for assets to identify services they can offer to 3rd parties, pricing model and how this will fit into a FI-WARE business frameworks.	All Assets
AP-4	Carlos to upload/forward the requirements discussed in this meeting. Check out other things needed and mail within the WP how the assets have to go forward (if there is a need for that). Questions we have for the testbed: Timeline, How the access will be.	Telefónica
AP-5	Architecture discussion: - All assets to analyze all possible common and specific functions for the REST APIs and provide which ones will be implemented.	All Assets.
AP-6	Architecture discussion: - Clarify “Distributed database” (Tali input about one UC project ticket) and its relation with the “middleware” discussion.	Telefonica
AP-7	Architecture discussion: - Dummy Use-cases. All team to progress the work in the mailing list and conference calls.	ALL
AP-8	CEP-PubSub Integration IBM to go back to Carlo and make a new suggestion with scenarios of integration. Workshop foreseen. Broader approach rather than just value-change event.	Tali, IBM
AP-8	NGSI topic - TI to make a proposal and discuss with the guys of IoT. Guy (IBM) to be involved in the discussions. Perhaps a Workshop within the chapter is needed.	Carlos, TI

AP-9	UC Project ticket backlog LS, Semantic App support, BigData assets done. For other assets: to be checked group-by-group in the call conference.	All
AP-10	D2.3. Overall deadline Feb 17th. We should have a draft for all by Feb 3rd. - To provide contributions. Use Query Broker as reference. Link for the notation of architecture figures: References: [1] - http://www.fmc-modeling.org/notation_reference [2] - http://www.yworks.com/en/products_yed_about.html Hi all, On Nov 30th, I sent a detailed email describing how I think we should approach development of the next major deliverable, i.e., the deliverable on FI-WARE Architecture Specifications. I hope that all of you had already read carefully its contents so that we all are on the same page. There were two questions that were pending to solve before starting the actual writing of the deliverable on the Wiki. I elaborate on the proposed approach for both of them here. 1. Conventions for figures to include in Architecture Specifications In order to get an harmonized set of specifications a convention for figures describing aspects of the Functional Architecture linked to a GE has to be adopted. After some analysis, Thomas and me have decided to propose following FMC conventions for "Block	All assets

Diagrams - Compositional Structures" defined in [1]. Note that the adoption of FMC conventions is limited to these Block Diagrams. Rest is not mandatory at all (nor indeed needed for this deliverable).

Unless we hear about any objection, this will be the adopted recommendation you should forward to your teams.

The gallery of basic elements used in FMC Block Diagrams is pretty simple, so I'm sure you could use any of your favorite editing tools for creating Architecture Description Diagrams, even powerpoint. For those who may not want to use powerpoint but a drawing tool that is better tailored to draw Diagrams, we may recommend yED (see [2])

2. Uploading contents on the Wiki

As already explained in one of my previous mails in response to a question made by Torsten, we will definitively go for developing contents of this deliverable on the Wiki. Therefore, one of the questions we should answer first is where on the Wiki.

Here you are a list of points that describe the approach I suggest we should follow:

- A new entry on the main home page of the FI-WARE Wiki will be created titled "FI-WARE Architecture". This will lead to a Wiki page where we will provide:
 - a short introduction of the goal of the Architecture Specifications
 - links to specific Wiki pages, one per chapter of FI-WARE. Each of these Wiki pages will be structured so that it includes an Introduction section (we will decide what comes here later) and a section per GE who should follow the structure we described in the email I sent on Nov 30 (note that the CORBA Event

	Service example I provided would map to the concept of a single GE in FI-WARE like the Pub/Sub Broker GE) ○ a link to a Wiki page to be titled "Bringing all pieces together", where we will elaborate on the description of how the different chapters will connect together from an architectural point of view, serving example (but generic) use case scenarios. ● We should include a section titled "Open Specifications" at each of the GE sections we have under the "Materializing the FI-WARE Vision" part of the Wiki. This section will contain a bullet list of two items, each one linking to a Wiki page. The first one will be titled "Architecture Specification" and the second one will be titled "Detailed Interface Specifications". The first one will be a direct link to the section dedicated to the GE under the "FI-WARE Architecture" part of the Wiki described in the point above (note this will allow to navigate to concrete GE Architecture Specifications from the "Materializing the FI-WARE Vision" part of the Wiki, but that is precisely something we want to achieve). The second one will be where the detailed description of what remains regarding the complete set of Open Specifications linked to a GE, that is, the detailed description (signature and behaviour description of provided APIs, definition of protocols, non-functional mandatory features, etc) of interfaces introduced in the Architecture Specification. A draft of the guidelines, based on the above descriptive text, will be made soon available on the FI-WARE Project Handbook available on the Wiki. Please don't hesitate to make any question or formulate	
--	--	--

	any doubt so that answers can help to enrich the guidelines. Please share this email with members of your team. Cheers, -- Juanjo	

7. Reference documentation

This minutes will be uploaded to Data/Context Docs folder in the FI-WARE forge as usual.