



eMontage: An Architecture for Rapid Integration of Situational Awareness Data at the Edge

Soumya Simanta
Gene Cahill
Ed Morris





Motivation – Situational Awareness

First responders and others operating in the “last mile” of crisis and hostile environments are already making use of handheld mobile devices in the field to support their missions.

Rapid Incorporation of New Data Sources

- Many data sources (real-time, historical, ...)
- Data is fragmented across different apps on the mobile device

Minimized Information Overload

- Edge users are under high cognitive load
- Information required is a function of user’s context and therefore dynamic

Simple Use

- Users are under high stress
- Small screen devices

Resource Constrained Hostile Environment





Hostile Environments Characteristics

Wimpy edge nodes

- Limited resources (CPU, battery and memory) on mobile nodes
- Example: Expensive computations on a smartphone may drain the battery fast

Limited or no end-to-end network connectivity

- Implicit assumption of WAN connectivity is not always valid
- Example: No access to internet during a disaster, DoS attack

High cognitive load

- Application latency and fidelity become important
- Example: A slow application will increase the cognitive load on the user

Bounded elasticity

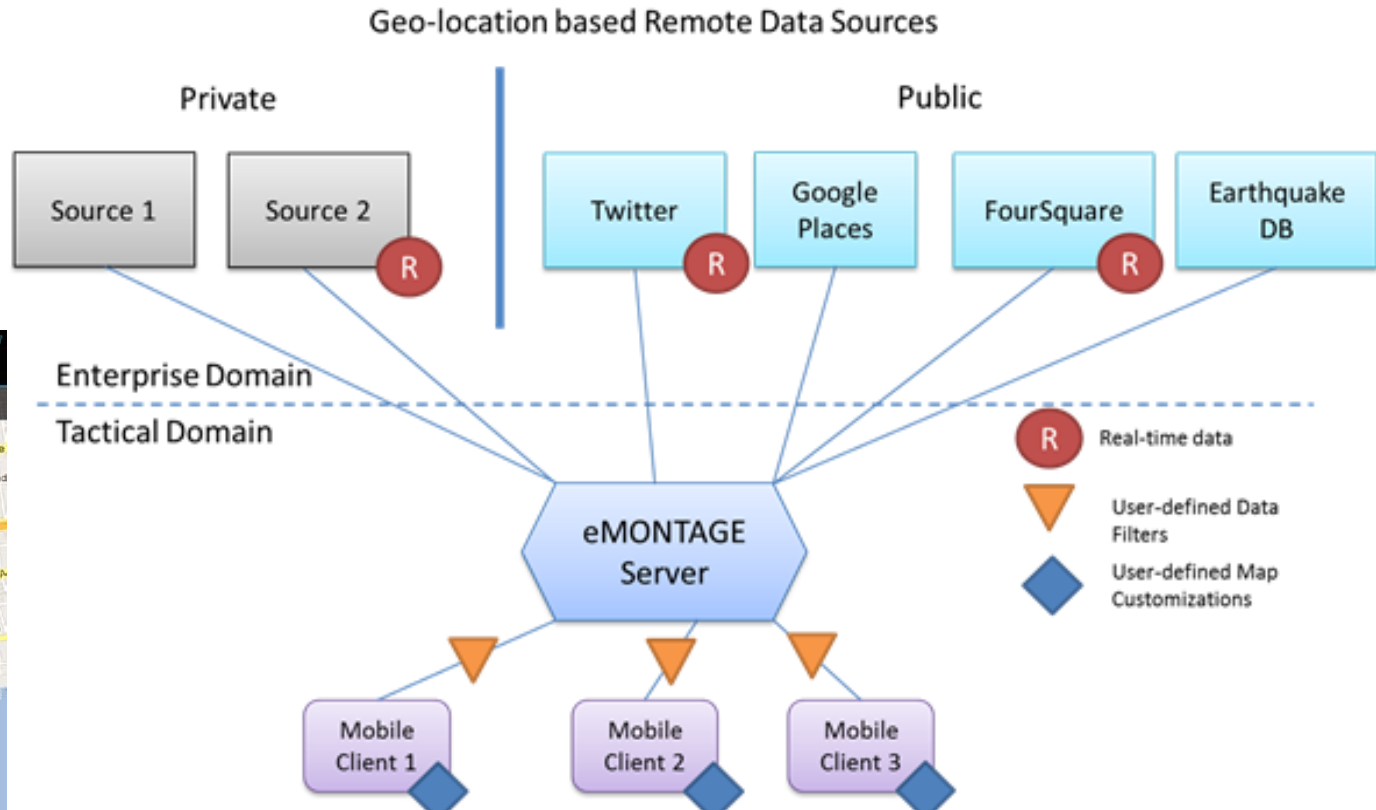
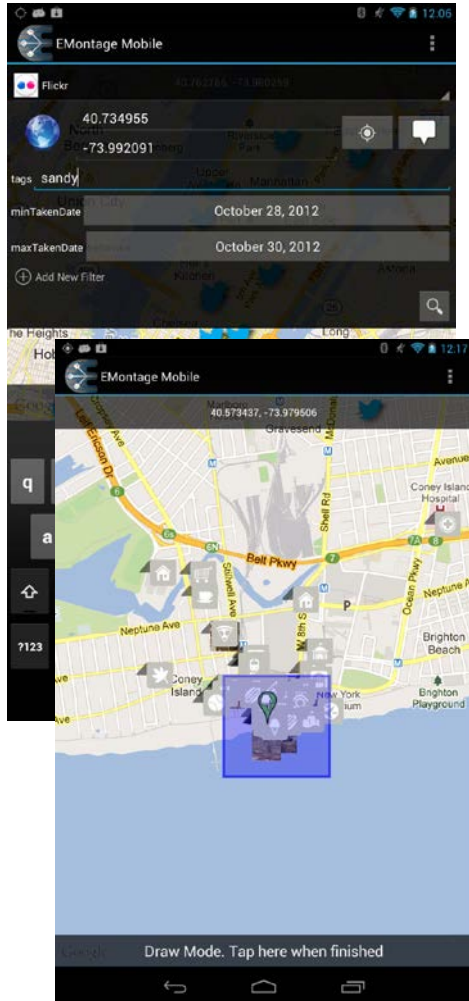
- Upper bound on number of consumers known in advance
- Example: Fixed number of first responders in a location

Dynamic environment

- Static deployment topologies cannot be assumed; Survivability essential
- Example: An automobile with a server may not be available



Context Diagram



Architecturally Significant Requirements

Extensibility

- Add new data source quickly with minimal impact on existing sources

Runtime Configurability

- Make data sources user configurable (e.g., using data filters) at runtime

Performance

- Minimize network bandwidth usage of the tactical network

Energy Efficiency

- Optimize energy consumption on mobile handheld devices

Usability

- Provide a responsive and unified user interface

Availability

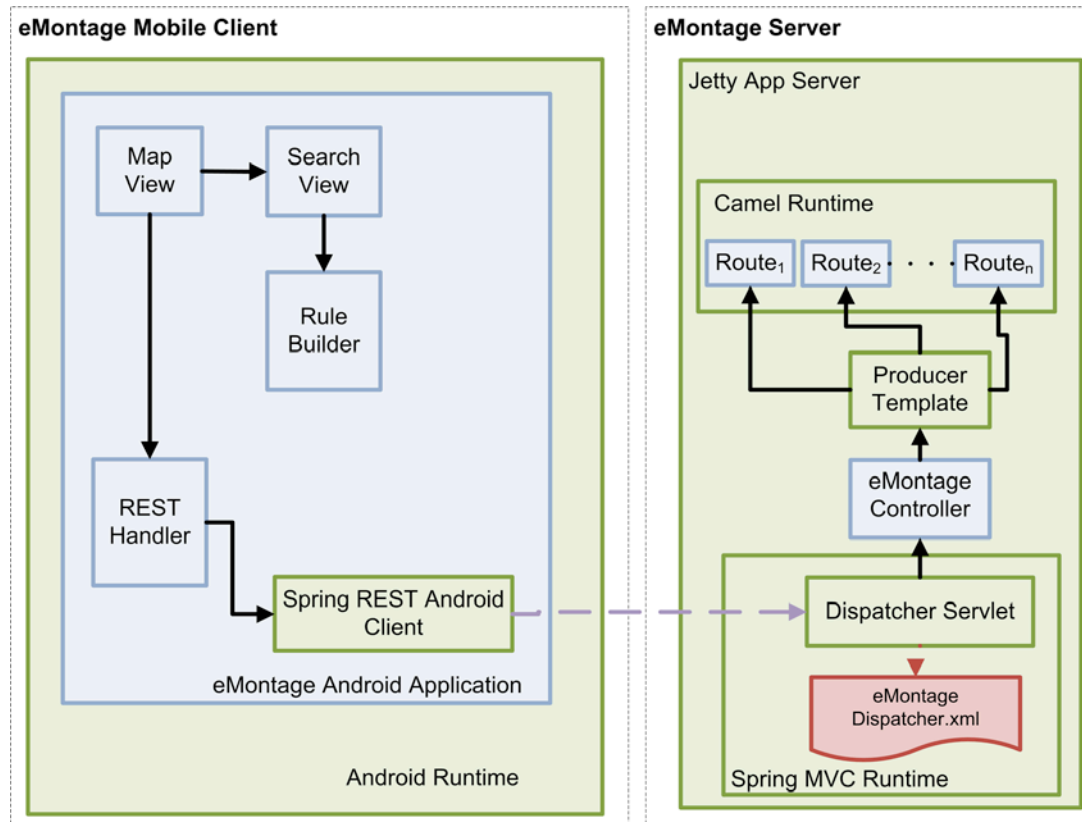
- Support intermediate disconnections with remote data sources

Security

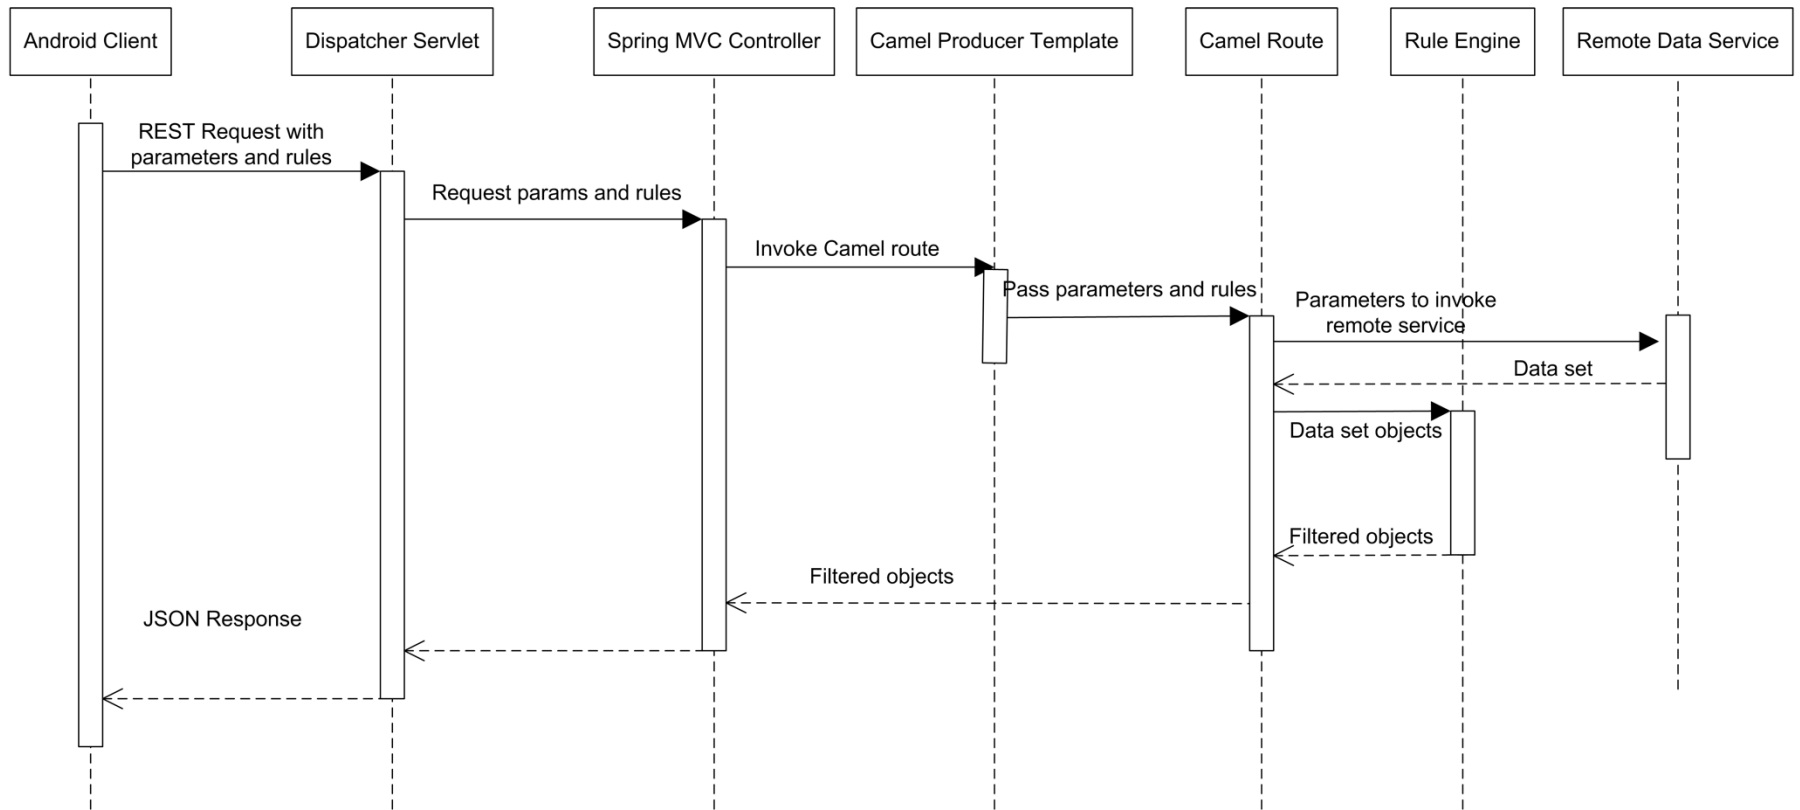
- Support existing security protocols and provide transport layer security



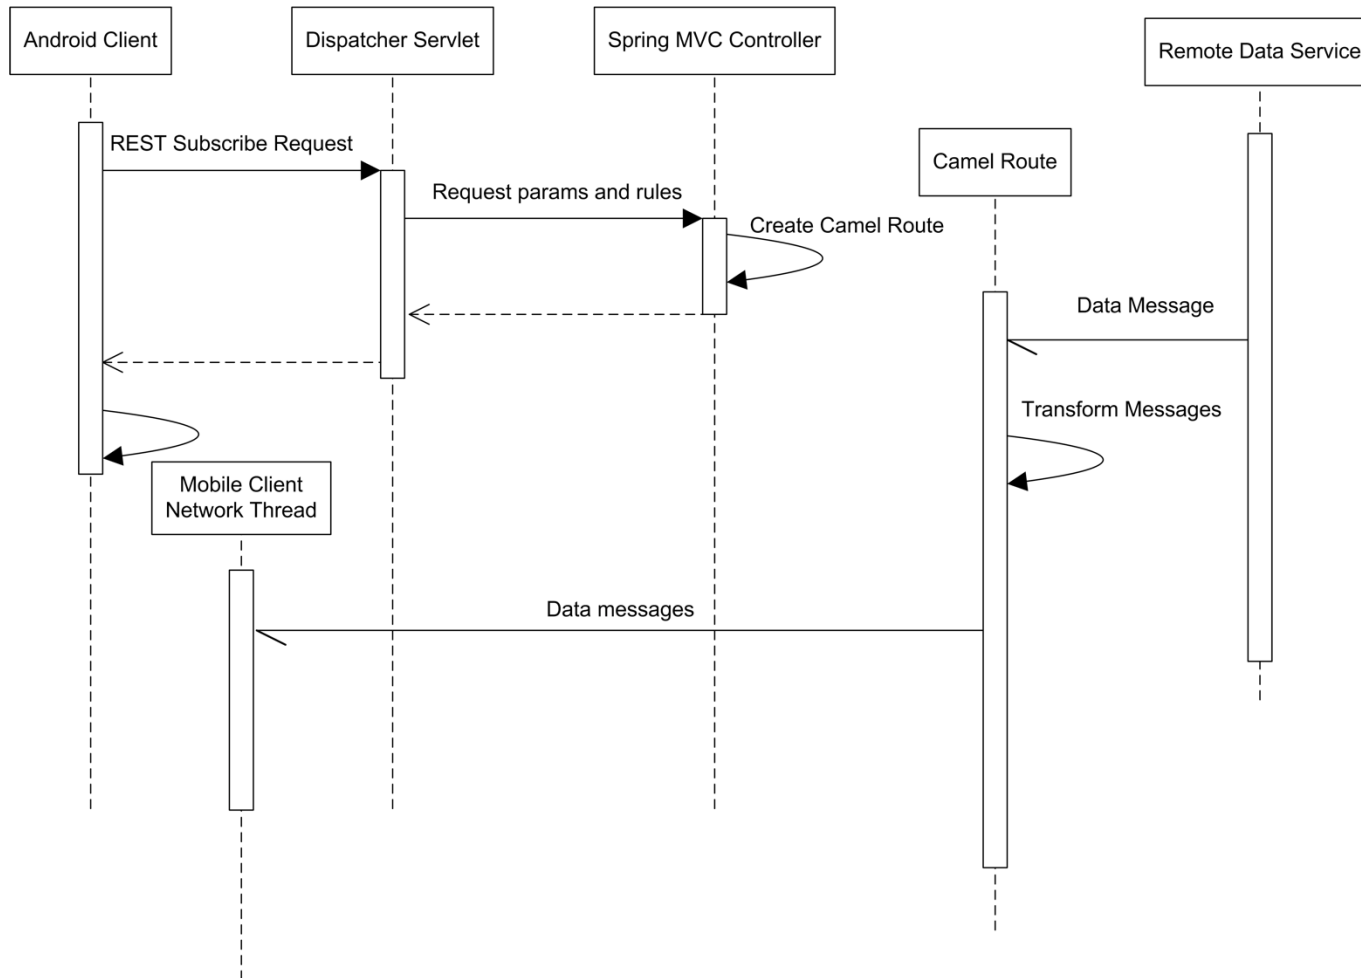
Runtime C&C View



Request Response Interaction




Publish Subscribe Interaction

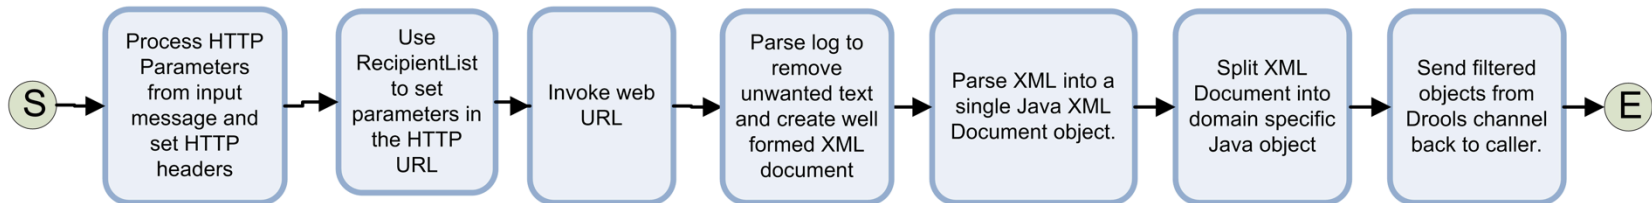


Example Routes

Google Place Route



National Weather Service Alerts Route



Extensibility – Adding New Data Sources

Problem

- New data sources are available in the field
- Adding and validating them is time consuming

Assumptions

- The data source has a remote API
- The data format is defined and stable

Solution

- Minimize coupling between data sources by encapsulating each data source
- Implement common connectors (request/response and publish subscribe)

Future extensions

Automate common tactical integration patterns to provide an end-user programming interface



Data Model

Data model is represented as objects (POJOs) shared between clients and server

Data model changes must be synchronized between clients and servers

- Our assumption: data model is “relatively” stable

Data model can be created in the following ways

- **Manual definition** – works best in case of a simple data model
- **Code generation** – WSDL2Java to generate code
- **Reuse existing library** – Twitter4J is an existing Java implementation of Twitter API



Mashup Mechanism

Merging data across data models is a mechanism to relate data across models.

- Example: foreign keys in a relational database

In eMontage, we assume a large proportion of situational awareness data has some form of geo-location associated with it

- use geo-location as the common key

All data is currently mashed up on a map-based interface.

- as long as two data elements from different data sources are referenced by geo-location (latitude and longitude), they will always be displayed correctly on a map
- the actual relating of information will happen with the user



Example Data Sources

Data Source Name	Data Format	Wire Protocol	Security Mechanisms	Data Model Complexity	User Interface Complexity
Google Places	JSON	REST	Token-based	Low	Geo-points on map
Twitter	JSON	REST	Token-based	High	Geo-points on map
FourSquare	JSON	REST	Token-based	High	Geo-points on map
Private Data source 1 (real-time)	XML	UDP	No security mechanism	Low	Polygons, Geo-points on map
Private data source 2 (historical)	SOAP	HTTP	Custom	High	Geo-points on map
National Weather Service Alerts	Custom log files	HTTP	No security mechanism	Medium	Polygons



Configurability- User-defined Runtime Filtering

Problem	Information overload
Assumption	The user knows what information they need in a particular context (e.g., location, keywords, date ranges)
Solution	Provide mechanisms that allow users to reduce the volume of information <ul style="list-style-type: none">• using rule-based filtering at runtime
Future extensions	Provide “data discovery” mechanisms (e.g., visualizations, clusters, outliers) when the user does not know what information they need





Usability - Unified User Interface

Problem	Data is fragmented across multiple applications and databases
Assumption	<ul style="list-style-type: none">• Data is geo-coded• A unified view provides more value compared to isolated views of data from different sources
Solution	Mashup of geo-code data viewed on a map allows <i>visual</i> unification of data
Future extensions	<ul style="list-style-type: none">• Provide other non-map based visualizations• Provide data join mechanisms






Performance - Minimized Bandwidth Utilization

Problem	Bandwidth is a scare resource at the “last mile” of edge
Assumption	Possible to have an intermediary node in the network
Solution	Add an intermediary node <ul style="list-style-type: none">• Use filtering at the source (only send information is required by the mobile nodes)• Transform to a more bandwidth optimized format (e.g., XML to JSON/Protocol Buffer)
Future extensions	Use protocol transformation (use a SPDY instead of HTTP)





Power Consumption - Offloading Expensive Computation

Problem	Mobile nodes have limited resources (CPU, battery and memory)
Assumption	Possible to have an intermediary node in the network
Solution	Perform expensive computation (e.g., XML parsing, multiple network calls) on a proximate, relatively resource rich node
Future extensions	Use multi-node cloudlets to increase performance and fault tolerance



Availability - Disconnected Operations

Problem	Edge nodes may have to work in disconnected or semi-connected mode (from enterprise/TOC network)
Assumption	<ul style="list-style-type: none">• Possible to deploy a resource-rich node locally (e.g., on an automobile)• Real-time data is generated <i>locally</i>• Possible to know in advance what data will be required for a mission (e.g., maps by locations)
Solution	Localize and cache data sources on a cloudlet.
Future extensions	Use persistent distributed caching. Adaptive pre-fetching to support intermittent disconnections



Architectural Alternatives

Native Mobile Client Only


- A **native, mobile client** app directly connected to backend data sources

Mobile Browser-Server

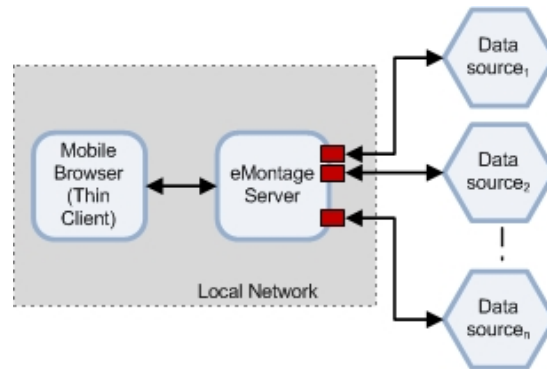
- A **mobile browser** client to a server that acts as an **intermediary** between the mobile client and the backend data sources

Native Mobile Client-Server


- A **native, mobile client** app connected to an **intermediary** server



Native Mobile Client Only



Mobile Browser-Server



Native Mobile Client-Server



Architectural Alternatives

	Native Mobile Client Only	Mobile Browser-Server	Native Mobile Client-Server
Reuse of COTS	Low	High	High
Protocol/Data format transformation	No	Yes	Maybe
Intermediate filtering	No	Yes	Yes
Bandwidth optimization	No	Yes	Maybe
Disconnected operations	No	Yes	Yes
Rich user interface	Yes	Yes	Maybe
Energy efficiency	No	Yes	Maybe
Fault Tolerance	High	Low	Low
Caching	No	Yes	Yes
Runtime modifiability	Low	High	High



Current and Future Work

More intuitive user interface

- Support other views of data
- Provide data exploration and discovery capabilities

Focus on performance

- Use caching
- Allow use of multiple processors/cores when possible

Add security mechanisms

- Use with Wave Relay radios
- Add transport and message level encryption

Integrate with edge analytics

- Build edge analytics techniques on top of current eMontage implementation



Questions



This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT. This material has been approved for public release and unlimited distribution except as restricted below.

This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

DM-0000349

