TEGRID: *Demonstration of a Semantic Web Environment*

Steven Gollery, Senior Software Engineer
Jens Pohl, Ph.D.**, Executive Director**

Collaborative Agent Design Research Center
Cal Poly, San Luis Obispo, California

Introduction

Over the past several years there has been an increasing recognition of the shortcomings of message-passing data-processing systems that compute data without understanding, and the vastly superior potential capabilities of information-centric systems that incorporate an internal information model with sufficient context to support a useful level of automatic reasoning.

The key difference between a data-processing and an information-centric environment is the ability to embed in the information-centric software some understanding of the information being processed. The term ***information-centric*** refers to the representation of information in the computer, not to the way it is actually stored in a digital machine. This notion of ***understanding*** can be achieved in software through the representational medium of an ontological framework of objects with characteristics and interrelationships (i.e., an internal information model). How these objects, characteristics and relationships are actually stored at the lowest level of bits in the computer is immaterial to the ability of the computer to undertake reasoning tasks. The conversion of these bits into data and the transformation of data into information, knowledge and context takes place at higher levels, and is ultimately made possible by the skillful construction of a network of richly described objects and their relationships that represent those physical and conceptual aspects of the real world that the computer is required to reason about.

In a distributed environment such information-centric systems interoperate by exchanging ontology-based information instead of data expressed in standardized formats. The use of ontologies is designed to provide a context that enhances the ability of the software to reason about information received from outside sources. In the past, approaches to inter-system communication have relied on agreements to use pre-defined formats for data representation. Each participant in the communication then implemented translation from the communication format to its own internal data or information model. While relatively simple to construct, this approach led to distributed systems that are brittle, static, and resistant to change.

It is the premise of the TEGRID (Taming the Electric Grid) proof-of-concept demonstration that, for large scale ontology-based systems to be practical, we must allow for dynamic ontology definitions instead of static, pre-defined standards. The need for ontology models that can change after deployment can be most clearly seen when we consider providing information on the World Wide Web as a set of web services augmented with ontologies. In that case, we need to allow client programs to discover the ontologies of services at run-time, enabling opportunistic access to remote information. As clients incorporate new ontologies into their own internal information models, the clients build context that enables them to reason on the information they receive from other systems. The flexible information model of such systems allows them to evolve over time as new information needs and new information sources are found.

The TEGRID Demonstration Context

Since mid-2001 the Emergency Operations Bureau of the Los Angeles Sheriff's Department has been assigned the additional task of coordinating the response to expected rolling electric power blackouts, as California's demand for electric power came perilously close to exceeding availability. While both the power outage areas and individual blackout periods are predefined in terms of a large number of power grid units that are distributed throughout the Los Angeles County, the emergency events that are likely to be triggered by blackout conditions (e.g., multi-vehicle accidents, carbon monoxide poisoning in enclosed parking garages, fires, criminal activities, and other disturbances) are less determinate.

The TEGRID proof-of-concept system has been designed to assist the Los Angeles Sheriff's Department by addressing this potentially chaotic situation in an autonomously evolving, just-in-time manner. TEGRID does not exist as a pre-configured system of tightly bound components that know about the existence of each other, have predefined connections, and predetermined capabilities. In fact at the beginning of the demonstration TEGRID, as a system, does not really exist at all. What does exist is a set of cooperating Semantic Web Services, based on standard Web Service specifications (e.g., SOAP, UDDI, WSDL, and XML) enhanced by the ability of sharing semantic-level descriptions of their own internal information models.

In essence TEGRID involves sharing information among a number of separate organizations, including local police stations, the Emergency Operations Bureau, a power supply management and monitoring organization, and a traffic control system. The proof-of-concept relies on a set of assumptions about the existing resources available from each of the organizations involved.

1. That each local sheriff's station has a database that includes (at least): current officer assignments; equipment manifests and status; and, priority infrastructure and intersections.
2. That the Emergency Operations Bureau has a list of Rapid Response Teams and their primary and alternative assignments.
3. That there exists some kind of Power Supply Organization that has a database of recent history of power consumption, plus the ability to provide a real-time feed of current power levels.
4. That there exists some kind of Traffic Control Organization that has some method of determining acceptable alternative routes for reaching a particular destination from a given starting location.

Another underlying assumption is that all of these organizations have Internet connections and either have an existing web site or are willing to establish one. TEGRID builds on these existing information and data sources to construct a web service infrastructure that allows information-sharing and automated decision-support.

Since the proof-of-concept system does not have access to live databases, it simulates them, using sample data to implement the demonstration scenario. There are also some potential applications that must exist in order to support the scenario, but are not part of TEGRID itself. For example, there is a requirement that new incidents (e.g., traffic accidents) would be reported to the local sheriff's stations before they are able to propagate through the system. Such a reporting application is assumed to exist, and has been simulated in order to produce the dynamic behavior called for in the demonstration scenario.

TEGRID features several kinds of web service providers. Each of these implements a set of operations that allows exchange of the information that makes the functioning of the system possible. These operations such as subscription, information transfer, warning and alert generation, discovery, and assignment, are the minimum necessary to provide the functionality described in the demonstration. More operations can be easily added as TEGRID's capabilities increase in the future.

In addition, TEGRID includes software agents with automatic reasoning capabilities. Some of these agents could conceptually be seen as services. For instance, the Station Monitor Agent is able to publish alerts that the local stations can subscribe to, and at the same time the Station Monitor Agent is able to subscribe to notifications of planned power outages. The relationship between agents and services is perhaps a fertile field for further investigation: When is it more useful to implement functionality as an agent, and when as a service? Are the two orthogonal? Is it reasonable to think that the same set of functions might be an agent from one point of view, but a service from another? Does an agent consume services, provide services, or both? Since it seems likely that the answers to these questions depend on the nature of the individual agent, the definition of a conceptual framework for making such determinations might be a productive future goal.

The Fundamental Web Service Elements

Within the Internet context of web services, TEGRID builds on a number of standard protocols and elements. These elements are combined into an executing software entity, capable of seeking and discovering existing web services, extending its own information model through the information model of any discovered web service, and automatically reasoning about the state of its internal information model. As shown in Fig.1, this entity or Cyber-Spider consists of three principal components: a web server; a semantic web service; and, an information-centric application.

The web server, utilizing standard Hypertext Transfer Protocol (HTTP), serves as the gateway through which the Cyber-Spider gains access to other existing web services. Web servers primarily provide access to Hypertext Markup Language (HTML) data sources and perform only simple operations that enable access to externally programmed functionality. However, these simple operations currently form the building blocks of the World Wide Web.

The second component of a Cyber-Spider is a semantic web service (i.e., a web service with an internal information model). A web service is accessed through a web server utilizing standard protocols (e.g., UDDI, SOAP, WSDL, SML) and is capable of providing programmed functionality. However, clients to a standard web service are usually restricted to those services that implement specific predefined interfaces. The implementation of web services in the Internet environment allows organizations to provide access to applications that accept and return complex objects. Web service standards also include a limited form of registration and discovery, which provide the ability to 'advertise' a set of services in such a way that prospective client programs can find services that meet their needs. The addition of an internal information model in a semantic web service allows the storage of semantic level descriptions (i.e., information) and the performance of limited operations on these semantic descriptions. In other words, the semantic web server component of a Cyber-Spider is capable of reasoning.

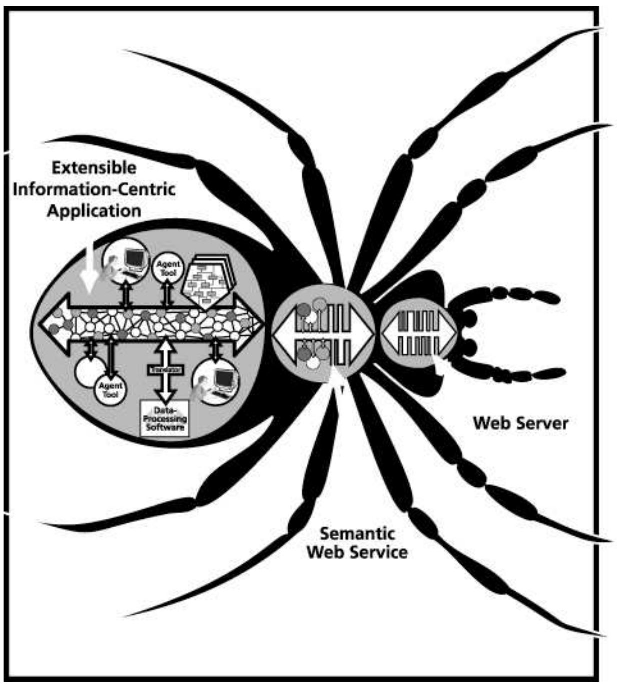

Fig.1: Anatomy of a Cyber-Spider

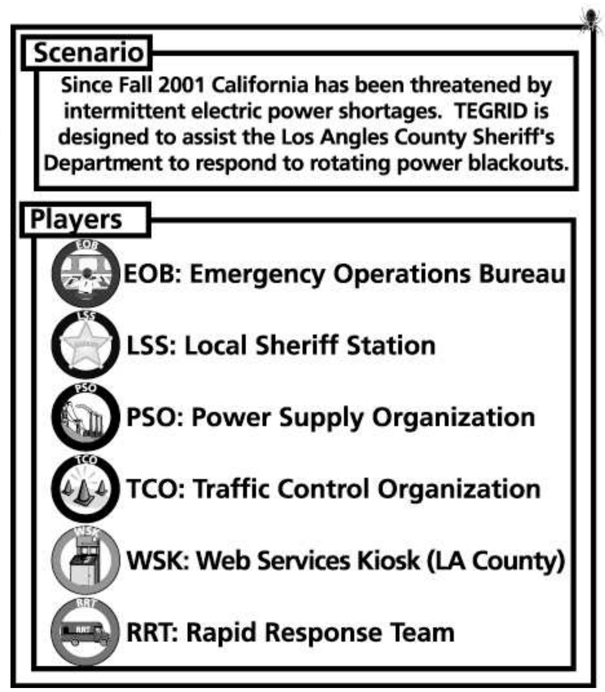

Fig.2: Cast of TEGRID players

The third component of a Cyber-Spider is one or more information-centric applications. These applications are designed to take advantage of the resources provided by a number of semantic web services, enabling them to reason about the usefulness of each service and support more sophisticated discovery strategies. Moreover, the application component is able to construct relationships among the information models of different services, with the ability to integrate services without requiring agreement on a common information model.

With these three components Cyber-Spiders are at least minimally equipped to operate in an Internet environment as autonomous software entities, capable of: discovering needed services; accepting services from external offerers; providing services to external requesters; gaining context through an internal information model; automatically reasoning about available information; extending their information model during execution; extending their service capabilities during execution; and, learning from their collaborations.

The TEGRID Players

The cast of players in the current TEGRID proof-of-concept demonstration includes six players or existing web services (Fig.2): the Emergency Operations Bureau (EOB) of the Los Angeles Sheriff's Department; several Local Sheriff Stations (LSS); a Power Supply Organization (PSO); a Traffic Control Organization (TCO); several Rapid Response Teams (RRT); and, a Los Angeles County Web Services Kiosk (WSK).

Fundamental to each player are three notions. First, each player operates as an *autonomous* entity within an environment of other players. Most, but not all of the other players are also autonomous. This requires the autonomous players to be able to discover the capabilities of other players. Second, each autonomous player has a sense of *intent* to accomplish one or more objectives. Such objectives may range from the desire to achieve a goal (e.g., maintain situation awareness, coordinate the response to a time critical situation, or undertake a predetermined course of action following the occurrence of a particular event) to the willingness to provide one or more services to other players. Third, each player (whether autonomous or not) is willing to at least *cooperate* with the other players. In some cases the level of cooperation will extend to a collaborative partnership in which the partnering players contribute to the accomplishment of a common objective. In other cases the cooperation may be limited to one player providing a service to another player, without any understanding or interest in the reason for the service request.

To operate successfully in such an autonomous Internet-based environment a Cyber-Spider player should be endowed with the following capabilities:

1. Subscribe to information from external sources
   (e.g., alerts, ontology extensions).
2. Accept subscriptions from external clients.
3. Dynamically change its subscription profile.
4. Extend its internal information representation.
5. Extend its own service capabilities.
6. Generate new agents for its own use.
7. Describe its own service capabilities to external clients.
8. Seek, evaluate and utilize services offered by external clients.
9. Provide services to external clients.
10. Describe its own (intent) nature to external clients.

The Cyber-Spiders in TEGRID are currently capable of demonstrating eight of these ten desirable capabilities. The ability of a Cyber-Spider to dynamically change its subscription profile, while technically a fairly simple matter, has not been implemented because it is not used in the demonstration scenario. The ability of a Cyber-Spider to describe its own nature to external clients, on the other hand, is technically a much more difficult proposition. It will require a Cyber-Spider to have an understanding of its personality as a collective product of its internal information model and the relationship of that model with the external world. At best this must be considered a challenging research area that is beyond the current capabilities of information-centric software systems.

The TEGRID Agents

Most of the reasoning capabilities available in TEGRID are performed by software agents that are components of the players (e.g., Cyber-Spiders). In other words, agents are predefined clients within player systems (i.e., information-centric applications) and perform internal functions that are necessary for the particular player to deliver its services and/or accomplish its intent. The following agents (i.e., collaborative tools) are available in the current TEGRID implementation:

Name of Agent	Owner	Description of Agent Capabilities
Risk Agent	EOB	Identifies high risk entities in the jurisdictional region of an activated LSS.

Deployment Agent	EOB	Determines whether RRT support is required for a particular activated LSS.
Power Level Agent	PSO	Determines if electric power demand has exceeded supply.
Situation Agent	EOB	Prepares and updates the 'EOB Situation Status Report'.
Station Monitor Agent	EOB	Identifies all LSSs that will experience power blackouts during the current and next blackout cycle.
Status Agent	LSS	Prepares and updates the 'LSS Situation Status Report'.
Local Station Agent	LSS	Determines whether sufficient local resources are available to deal with current conditions.
Scheduling Agent	EOB	Assigns RRTs and equipment to situations requiring RRT involvement.
Incident Agent	EOB	Monitors the response to a particular situation supported by one or more RRTs.
Routing Agent	TCO	Determines alternative routes to a particular situation location.

Demonstration Objectives

Stated succinctly, the objective of the TEGRID scenario is to demonstrate the discovery, extensibility, collaboration, automatic reasoning, and tool creation capabilities of a distributed, just-in-time, self-configuring, collaborative multi-agent system in which a number of loosely coupled Web Services associate opportunistically and cooperatively to collectively provide decision assistance in a crisis management situation. Specifically, these capabilities are defined as follows:

Discovery: Ability of an executing software entity to orient itself in a virtual cyberspace environment and discover other software services.

Extensibility: Ability of an executing software entity to extend its information model by gaining access to portions of the information model of another executing software entity.

Collaboration: Ability of several Web Services to collaboratively assist each other and human users during time critical decision making processes.

Reasoning: Ability of a software agent to automatically reason about events in near real time under time critical conditions.

Tool Creation: Ability of a Web Service to create an agent to perform specific situation monitoring and reporting functions.

Players' Intent

The TEGRID players or Cyber-Spiders are initialized with intent or willingness to cooperate based on their role and operational responsibilities, as follows:

EOB (Emergency Operations Bureau): To be immediately informed of imminent power blackout conditions, to coordinate all assistance to LSSs, to maintain situation awareness, and to take over local command responsibilities when conditions require actions that cross the jurisdictional boundaries of two or more LSSs.

LSS (Local Sheriff Station): To activate a predefined response plan as soon as it receives notification (from the EOB) that a power blackout condition is imminent within its jurisdiction, to respond to new emergency missions in its jurisdictional area, to provide RRTs to the EOB, and to request assistance from the EOB.

PSO (Power Supply Organization): To share information relating to the current status of power demand and availability with subscribers, to provide subscribers with information relating to a predefined rolling power blackout schedule on request, and to alert subscribers whenever the schedule is intended to be implemented.

TCO (Traffic Control Organization): To share information relating to historical traffic flows under typical conditions with subscribers, to provide subscribers with information relating to traffic control capabilities (e.g., types and location of traffic signals, sensors, and web-cameras), and to provide subscribers with alternate traffic routes on request.

RRT (Rapid Response Team): To share information relating to its current mission and location with subscribers, to execute missions requested by the EOB, and to provide assistance to any assigned LSS, and to request assistance from the EOB.

The TEGRID Demonstration Scenario

Armed with their individual intent and intrinsic Cyber-Spider capabilities (i.e., ability to: discover useful web services; subscribe to information and accept subscriptions from external clients; extend their internal information models; describe and provide services to external clients; seek, evaluate and utilize services offered by external clients; and, extend their own service capabilities by generating new agents) the players commence their partly intentional and mostly opportunistic interactions.

<u>Orientation</u>

The players orient themselves in the virtual cyberspace environment by accessing one or more directories of available services and registering an information subscription profile with those services that they believe to be related to their intent (Fig.3).

EOB (Emergency Operations Bureau): Accesses the WSK (Los Angeles County Web Services Kiosk) based on its predefined authorization level, and:

Subscribes to any service changes in the WSK.
Finds the PSO address which it was seeking.
Discovers the TCO.
Discovers all of the LSSs.

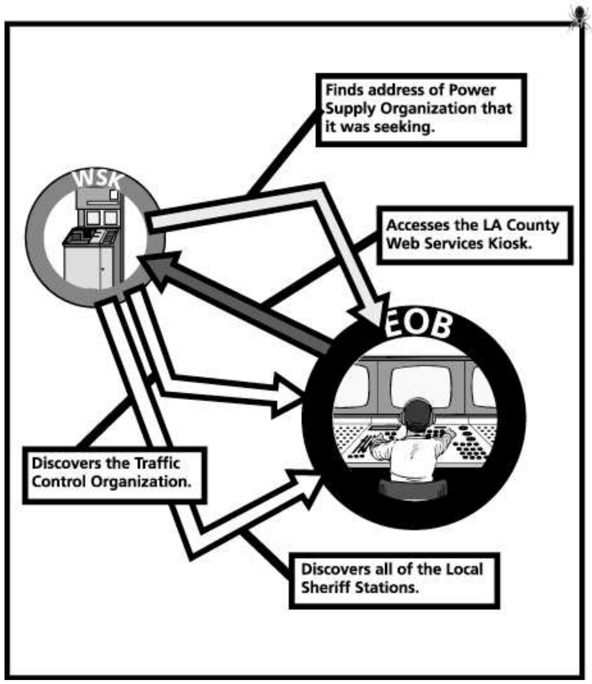

Fig.3: Orientation and discovery

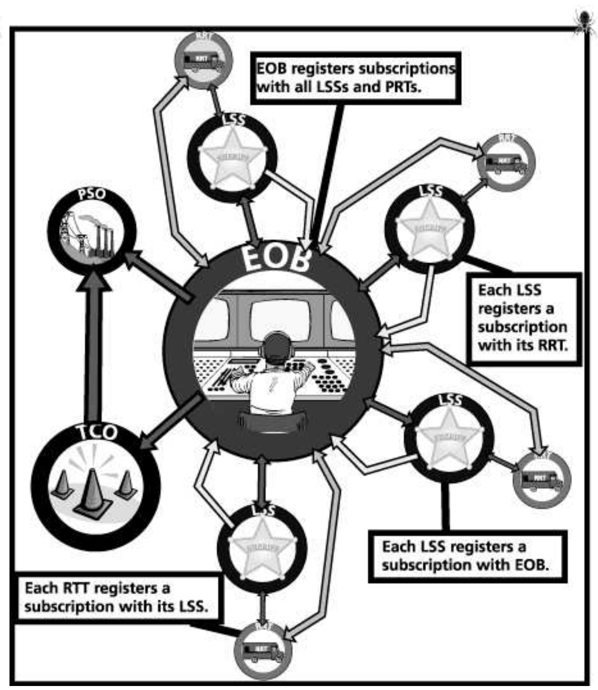

Fig.4: Information subscription

Subscription

The players access the services that they require to achieve their intent, register appropriate subscription profiles, and query for information that they believe to have a need for (Fig.4).

EOB (Emergency Operations Bureau): Registers a subscription profile with each LSS (Local Sheriff Station) that includes all current police unit locations, mission completion events, new mission events, and any information changes relating to the availability of its RRTs (Rapid Response Teams).

Queries each LSS (Local Sheriff Station) for all information relating to its RRTs (Rapid Response Teams) and extends its information model.

Registers a subscription profile with each RRT (Rapid Response Team) that includes its current location and mission.

Registers a subscription profile with the PSO (Power Supply Organization) that includes the current status of electric power demand and availability, and any change in its intention to implement the predefined rolling power blackout schedule.

Registers a subscription profile with the TCO (Traffic Control Organization) that includes any change in the status of traffic signals, sensors, and web-cameras.

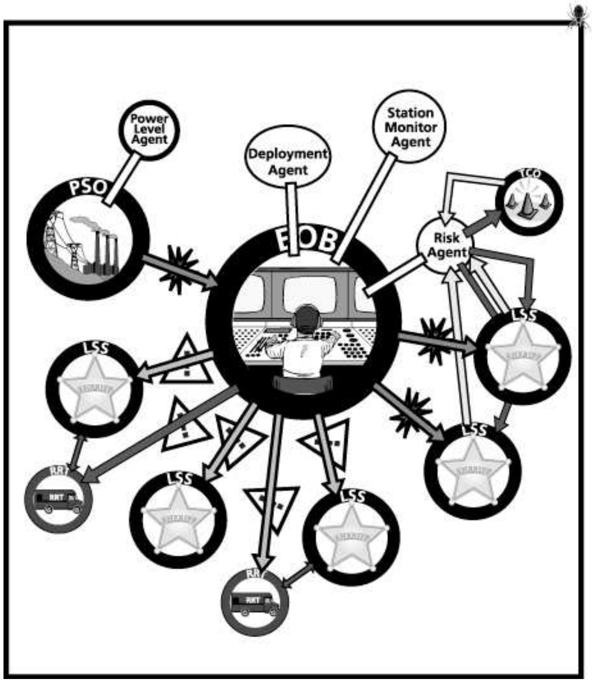

Fig.5: Power supply 'Warning'

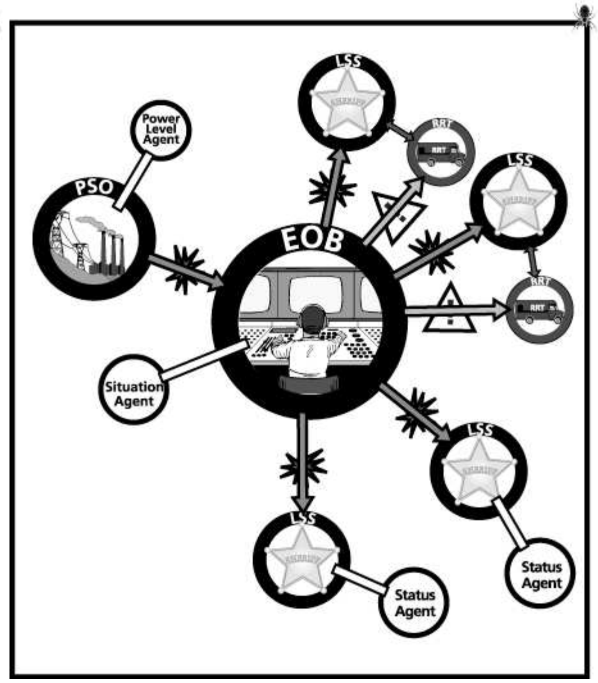

Fig.6: Power outage 'Alert'

LSS (Local Sheriff Station): Each LSS responds to the EOB (Emergency Operations Bureau) registration by registering a corresponding subscription profile with the EOB that includes the current mission and location of its RRTs (Rapid Response Teams), any EOB requests and orders to this LSS, and changes in the current 'situation status report' maintained by the EOB.

Each LSS (Local Sheriff Station) registers a subscription profile with its RRTs (Rapid Response Teams) that includes the current mission and location of the RRT, mission completion events, and new mission events (this duplication of its EOB (Emergency Operations Bureau) subscription profile allows the LSS to verify the accuracy of this portion of the 'situation status report' maintained by the EOB).

TCO (Traffic Control Organization): Registers a subscription profile with the PSO (Power Supply Organization) to include the location of all current power blackout areas.

RRT (Rapid Response Team): Registers a subscription profiles with the EOB (Emergency Operations Bureau) that includes any requests or orders to this particular RRT (Rapid Response Team), and any changes in conditions that impact the current mission and location of this RRT.

Registers a subscription profile with its home base LSS (Local Sheriff Station) that includes any request for information, and any 'situation status report' maintained by this LSS.

Power Outage Notification

The PSO (Power Supply Organization) alerts its subscribers that a rolling power blackout condition is imminent (i.e., will commence per predefined schedule within 15 minutes) (Fig.5).

PSO (Power Supply Organization): Utilizes its Power Level Agent to continuously monitor the relationship between power demand and supply. The PSO determines that demand is close to exceeding supply and sends an Alert to all appropriate subscribers.

EOB (Emergency Operations Bureau): Receives an Alert from the PSO (Power Supply Organization) that the predefined rolling power blackout schedule will be implemented within 15 minutes.

Utilizes its Station Monitor Agent to identify all LSSs (Local Sheriff Stations) that will experience power blackouts in their jurisdiction.

Warns all LSSs (Local Sheriff Stations) of imminent power blackout condition.

Alerts all LSSs (Local Sheriff Stations) in whose jurisdictions blackouts will occur and requests them to commence immediate implementation of their respective 'blackout response plans'.

Warns the RRTs (Rapid Response Teams) assigned to assist the LSSs (Local Sheriff Stations) in whose jurisdictions the first set of blackouts are scheduled to occur, to prepare for potential deployment.

Utilizes its Risk Agent to identify all high risk entities in the jurisdictions of the activated LSSs (Local Sheriff Stations). Utilizes its Deployment Agent to determine whether RRT (Rapid Response Team) involvement is anticipated under normal conditions.

LSS (Local Sheriff Station): Each LSS assumes 'alert' status. The LSSs in whose jurisdictions the first set of blackouts is scheduled to occur, prepare for deployment.

RRT (Rapid Response Team): The RRTs notified by the EOB (Emergency Operations Bureau) assume 'alert' status in preparation for potential deployment.

Power Outage Implementation

The PSO (Power Supply Organization) alerts its subscribers that the predefined rolling power blackout schedule has been implemented (Fig.6).

PSO (Power Supply Organization): Utilizes its Power Level Agent to determine that demand has exceeded the availability of electric power.

EOB (Emergency Operations Bureau): Receives an Alert from the PSO (Power Supply Organization) indicating that the predefined rolling power blackout schedule has been implemented.

Utilizes its Situation Agent to prepare the first version of the 'EOB Situation Status Report'.

Alerts all LSSs (Local Sheriff Stations) in whose jurisdictions the next scheduled set of blackouts will occur, to prepare for potential deployment.

Warns the RRTs (Rapid Response Teams) assigned to assist the LSSs (Local Sheriff Stations) in whose jurisdictions the next set of blackouts are scheduled to occur, to prepare for potential deployment.

LSS (Local Sheriff Station): All activated LSSs utilize their Status Agent to prepare the first version of their 'LSS Situation Status Report'.

The LSSs (Local Sheriff Stations) in whose jurisdictions the next set of blackouts is scheduled to occur, prepare for deployment.

Traffic Accident in Power Outage Area

A multi-car traffic accident occurs in a blackout area located within the jurisdiction of a particular LSS (Local Sheriff Station) (Fig.7).

EOB (Emergency Operations Bureau): Receives an Alert from a LSS (Local Sheriff Station) that a multi-car traffic accident has occurred on State Highway 5 south of Harbor Freeway.

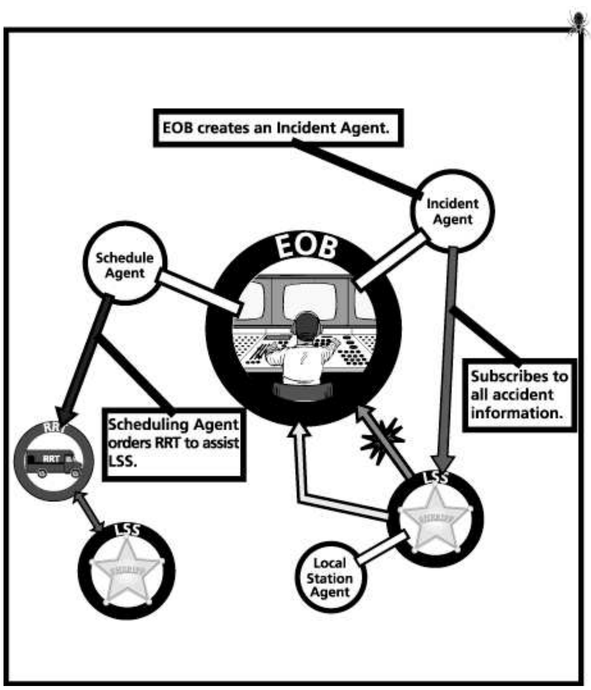

Fig.7: Traffic accident 'Alert'

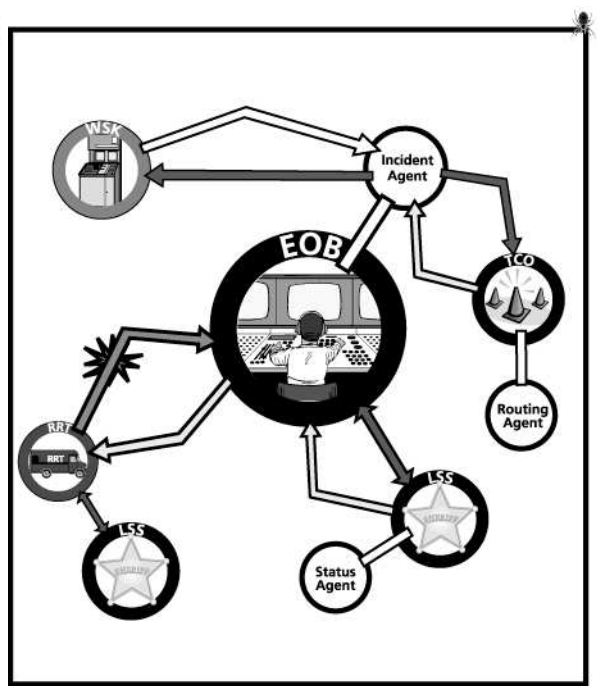

Fig.8: Routing assistance request

LSS (Local Sheriff Station): Utilizes its Local Station Agent to determine that it has insufficient resources to deal with the multi-car traffic accident.

EOB (Emergency Operations Bureau): Receives a request for assistance from the LSS (Local Sheriff Station) to deal with the multi-car traffic accident.

Utilizes its Scheduling Agent to assign a RRT (Rapid Response Team) and equipment to the multi-car traffic accident.

Creates an Incident Agent to monitor the response to the multi-car traffic accident.

The new Incident Agent subscribes to the LSS (Local Sheriff Station) in whose jurisdiction the multi-car traffic accident has occurred (to obtain all information about this accident from now on).

Routing Assistance Required

The dispatched RRT (Rapid Response Team) cannot reach the multi-car traffic accident due to traffic congestion and requests assistance in determining an alternative route (Fig.8) to the accident.

RRT (Rapid Response Team): Sends alert to the EOB (Emergency Operations Bureau) and requests assistance in determining an alternative route to the traffic accident.

EOB (Emergency Operations Bureau): Utilizes its Incident Agent to determine an alternative route. The Incident Agent accesses the WSK (Los Angeles County Web Services Kiosk) and discovers the TCO (Traffic Control Organization). It then registers a subscription profile with the TCO that includes routing information, and requests assistance in determining an alternative route to the traffic accident.

TCO (Traffic Control Organization): Receives the request for assistance from the EOB's (Emergency Operations Bureau) Incident Agent and utilizes its Routing Agent to determine an alternative route to the traffic accident.

Sends the alternate route to the EOB's Incident Agent..

EOB (Emergency Operations Bureau): Responds to the RRT (Rapid Response Team) by sending it the alternate route to the traffic accident.

Significance of the TEGRID Demonstration

The TEGRID proof-of-concept project was undertaken by the Collaborative Agent Design Research Center (CADRC) at Cal Poly (San Luis Obispo) as a small internally funded research endeavor with three objectives. The first objective was to explore the main capabilities that would be required of web service type entities (i.e., Cyber-Spiders) serving as largely autonomous decision-support components in a self-configuring, just-in-time, intelligent decision-assistance toolkit of collaborating software agents. Second, to determine if the currently available information-centric software technology could support at least basic (i.e., meaningful and useful) implementations of these required capabilities. And, third, to build a working experimental system that could serve as a test bed for longer term research studies focused on the behavioral characteristics of self-configuring intelligent systems in general, and the ability of such systems to deal with specific kinds of dynamic and complex problem situations.

The principal capabilities that are required by a Cyber-Spider to support the desired self-configuring, just-in-time, intelligent decision-support behavior have been identified and demonstrated in the TEGRID test bed environment, at least at a base level of functionality. These capabilities include the ability to: discover desired existing external services; accept and utilize services from external offerers; provide services to external requesters; gain understanding through the context provided by an internal information model; automatically reason about available information within the context of the internal information model; extend the internal information model during execution; spontaneously generate new agents during execution as the need for new capabilities arises; and, learn from the collaborations that occur within the cyberspace environment.