NuParadigm

10th International Command and Control Research and Technology Symposium
The Future of C2

Resolving the Problem of Aligning Communities of Interest, Data Format Differences, Orthogonal Sensor Views, Intermittency, and Security – DoD Homeland Security Command and Control Advanced Concept Technology Demonstration

Coalition Interoperability

Harry R. Haury
CEO
NuParadigm Government Systems
16091 Swingley Ridge Rd., Suite 160
Town and Country, MO 63017
636-537-5558 ext.-226 office
636-537-5262 fax
hhaury@nuparadigm.com

Introduction

The existence of diverse sensor, human intel and assessment information within isolated communities of interest creates blind spots in classic command and control environments where it is too expensive to instantiate standing capabilities to address a particular threat or align organizations to asymmetric attacks prior to an event. The problem is further exacerbated by the relatively rigid definition of community boundaries and views of relevant data within a community. The DoD/OSD Homeland Security Command and Control ACTD has attempted to address specific issues surrounding the problem of sharing C2 data between organizations and realigning data to address different roles and missions that might consume such data. In order to meet Department of Defense objectives to extend reach and improve agility through net-centricity, the ACTD implemented a sophisticated alerting framework to integrate various data sources and direct appropriately filtered and formatted data to users that have authority and need to see it. The design of the system, known as the Homeland Security On-Line Services Alerting Framework addressed the very real need to dynamically realign mission threads, organizations, and C2 systems to address the creation of new communities to address highly eccentric asymmetric attacks. Such dynamic capabilities also infer the use of secure service oriented applications that eliminate the provisioning problems associated with traditional C2 systems. Further the reach to civilian participants with a minimum of advanced training requires the implementation of highly automated simple to use systems, to this end the Alerting Framework has implemented a flexible interface for remote subscription and user characterization that directly modifies rules based data distribution driven from the edge.

This paper will address the implementation of the framework, the architectural approach to designing the system, the mission the system was designed to address, the technical and organizational approach to implementing the information sharing architecture, and the resultant capability delivered by the program.

The Problem and Solution Requirements

The implementation of GIG-BE, JTRS, NCES, GCCS/FCS, WIN-T and other advanced systems designed to support the deployment of network-centric warfighting will have a critical impact on requirements for data transmission, information assurance (“IA”), ad-hoc associations, support for communities of interest, and cross-boundary integration. These issues will be further complicated by the asymmetry of the connected networks and the extraordinary number of deployed users, sensors and systems that will come together. There is a dangerous tendency toward over simplification of the requirements that net-centricity

will place on the networks-----such deployments will require much more than simply adding bandwidth, transport encryption and IPv6.

The changes in application architecture wrought by network-centricity will dramatically increase traffic across certain network segments. Improper design and inefficient deployment of services will exacerbate the problem and multiply bandwidth requirements by as much as several orders of magnitude on parts of the network. Further, trust relationships between systems, users, data and roles will have to be managed in a standard, easily accessible fashion to allow inter-operability, effective reuse and secure operation without inappropriate limitations on functional capabilities. The technology for managing these ad-hoc trust relationships has yet to be developed and accredited.

This shift of paradigms significantly alters the mission of the network designer when trying to anticipate the underlying approach to delivering effective applications across the networks. Hundreds of specific architectural issues arise, but they can be broadly categorized into two categories:

Comparable Data Models – Net-centricity and the emphasis on interoperability are unlike any previous military application deployment model. The Internet and the broad interaction of commercial and private enclaves across the Internet better reflect the design and flexibility requirements for appropriate network infrastructure than anything currently deployed by the military for tactical and operational support. There are a host of obstacles to effective deployment of large-scale service-based applications.

Many of these problems have been solved architecturally within the commercial community and, to this end; much can be learned from the evolving architecture of the commercial Internet to identify the network requirements for supporting network-centricity and services-oriented architectures.

Outside Current System Capabilities - With all the similarities to the Worldwide Web, the military's requirements still go far beyond the current functional capabilities of the Internet. Specific functional extensions will be required that do not exist in the commercial network infrastructure. If these new architectural requirements are not anticipated properly, the promise of net-centric warfare cannot be achieved.

To understand the challenges faced by designers of these advanced networks, one needs to understand the fundamental differences between the broad deployment of service-based applications, as anticipated by net-centricity, and traditional silo-oriented systems that share network bandwidth, but have limited deployment reach. Net centricity changes the basic concept of how and where application functions are deployed. As such, these systems require new views of **many** common tenets in application implementation; some of the prominent factors that must be considered include the impact of:

a) ***Service Oriented Architectures –*** There are literally hundreds of important considerations in deploying service-based applications ranging from obvious to obscure, however, some of the more important architectural concepts need to be understood in order to begin to appreciate how system capabilities and performance are significantly impacted by the interaction of the following relatively simple issues.

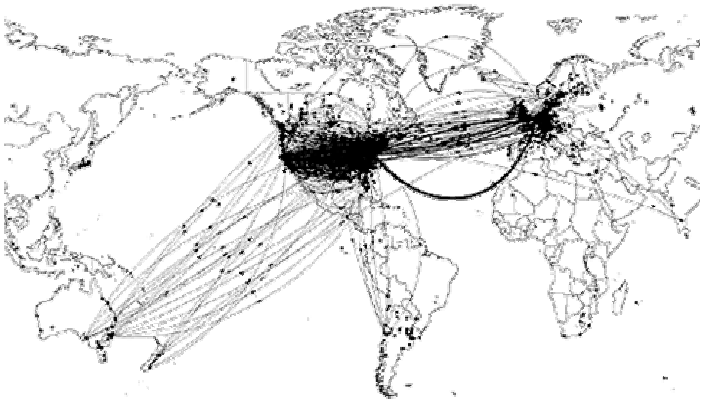

1. "Connectionless" interactions allow arbitrary dynamic movement of data between systems that merely have each other's address and understand the protocol for communication. Adherence to standards enables users, data, and systems to interact in an ad-hoc way without needing to be aware of each other in any specific fashion beforehand. Further, the use of standards-based processes enables such connections to be made without special or proprietary software. The problems come from the need to redirect traffic, maintain context between systems/steps in a process, implement a rational security model and provide session management. All of these requirements are complicated by the connectionless asynchronous architecture.

2. Much of the strength of the Worldwide Web comes from the virtual reference to resources and services through network services such as static/dynamic

DNS, instant messaging, and various connection proxies. This allows the network infrastructure itself to forge and manage the connections between systems, rather than having to manage the logic within the applications. This type of environment also supports redirecting connections away from an unavailable or dysfunctional resource to one that is available. The virtualization of routing is a requirement for any service-based architecture if functionality and capacity are to be effectively managed. Further, this virtual referencing needs to be able to route to a community or subscription list in order to manage multipoint connectivity. The more robust the proxy capability between systems, the easier it is to connect and manage them.

3. Much of the reliability, flexibility and scalability of web-based applications comes from the multi-tiered separation of data sources, application logic, presentation logic and generic network services. This enables various capabilities to be brought together in new and unique ways; it also allows portions of a system to be modified or improved without impacting the other parts of an application. This separation or "encapsulation" supports massive scaling through the seamless instantiation of a particular process on multiple machines and in multiple physical locations. However, this implies another significant layer of complexity regarding IA; the commercial internet has for the most part dispensed with this requirement except in robust B-to-B applications. The military does not have this option and will have to develop a model from scratch by leveraging and extending what is already known from the commercial sector and current deployments.

4. The migration of applications toward standards-based connections supports the reuse of code from one application to the next. The explosion of applications and functionality across the web has resulted directly from the standardization of protocols for communicating between systems. To the degree that a particular standard or subset thereof can be established and reused, the more leverage that will be gained. It is important to understand that there is a dark side, the broader use of standards that do not have an effective IA framework results in networks that are increasingly susceptible to hacking, manipulation and destruction. This effect is easily observed on the Internet today; examples include the number of attacks on Microsoft Outlook as compared to proprietary email systems. The better known the platform and

the more broadly implemented, the more important the proper design of IA becomes and the more likely it is to be attacked.

5. The task of providing seamless connectivity is one of the most important aspects of the Worldwide Web. Based on accepted standards ranging from the transport layer through the application layer, this seamless connectivity and interoperability make it extremely simple to deploy applications. By plugging a machine or process into this standards framework, it becomes extraordinarily simple to interface components together. This standardization is much more than packet structure; it includes all types of messaging protocols for http, ftp, JMS, .NET and many others.

6. The Internet has led to the evolution of applications being deployed across the network. As the understanding of this new abstraction has grown, it has become apparent that the traditional concepts of 1) what belongs to the application, 2) what belongs to the Operating System, and 3) what belongs to the network, have begun to change rapidly. Examples of this new approach abound, but a few of the more common are DNS, Dynamic DNS, third party certificate validation, instant messaging, ad-hoc peering through IP and identity correlation.

7. Service and source proxies are necessary for protection and network management. Proxies allow load management, redirection to proximate resources, and resilient connection redirection. Proxies in general change a reference from one virtual resource to another; in this way the application does not have to perform such resource arbitration and, if efficiently implemented, the proxies can also handle silent format transformation.

8. Traditional router-based IP protection and service side firewalls are insufficient to protect from denial of service attack. There are a number of reasons for this, but the primary one is that the routing network in most systems is extremely permissive if traffic is not blocked before it hits the network. To achieve an acceptable level of security, it is important to prevent mal-formed traffic or non-permitted traffic from reaching the network. Further, if it does penetrate the network boundary, the network itself needs to be capable of shedding it. There are a number of practical methods for doing

this, but they have not gained wide acceptance across the Internet except for high-risk transaction networks. Network-centric warfare will require a much greater investment in protection from traffic attacks than the generic Internet. Proxies traditionally provide protection by hiding information about what sits behind them from the outside world. At an object-level they can function as an "information guard", testing object specific IA and validating message permissions before passing an object through the network.

9. Traditional modeling and sizing of networks using IP packet counts is inadequate. There are a number of factors affecting total effective throughput beyond raw message size, these include, among other things:

    - the impact of the content or message propagation models,
    - content caching,
    - asynchronous queuing,
    - the design of the information assurance model,
    - methods for synchronization,
    - traffic management techniques,
    - bandwidth negotiation,
    - QoS assertion,
    - bandwidth reservation techniques,
    - CoS (Class of Service) negotiation,
    - content dissemination models, and
    - message recovery techniques.

    Effective available network capacity is heavily impacted by the availability and use of appropriate transaction management capabilities. These effects have often been ignored in simpler systems; however, in the applications contemplated by network-centricity, they could easily have a one to three order of magnitude impact on traffic across any particular network segment and will have to be anticipated. These effects are discussed in more detail below.

Policy or Role-based connectivity – Net-centricity implies the ad-hoc policy-based access of many users and systems to numerous applications and datasets that are distributed across the networks. The scope and scale of the envisioned interaction is unlike

anything that has preceded it. Systems of systems coming together with participant populations outside any particular application's administrative boundaries will create an immensely complex and difficult environment to deploy, manage, and maintain. Many of the difficulties arise from the architectural problem generally known as the n-squared problem, which must be ameliorated if the GIG vision is to become a reality. Standards-based open connections across standard networks substantially mitigate the difficulties of the n-square problem, but several factors, including: overlays on the network such as HAIPE-like virtual end-to-end piped connectivity; fixed physical connection references; proprietary connections; deployment of services with a single physical instantiation; etc. serve to complicate the environment. Further, the use of classic user access control and administrative processes make n-square mitigation nearly impossible. All participants in the network will require characterization and identity validation across security boundaries; this applies to users, services, networks, applications and data.

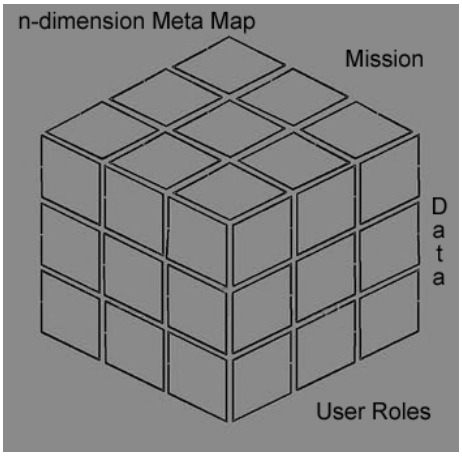

Ambiguity of Associations or Community Definitions – Ambiguity is the enemy; the network will intelligently bring together users and resources. The problems are directly analogous to the difficulties making a search using a web-search engine. If you cannot precisely and uniquely describe what you want, you will get thousands, perhaps millions of unwanted responses. Further, if other people do not use similar descriptive language to characterize their information, you will miss important data. Metadata used to associate roles, content, urgency, relevance, security and so on will determine whether the net-centric concepts of communities of interest will work. It will be just as dangerous to be inundated with data noise or allow improper data access, as it will be to miss important information.

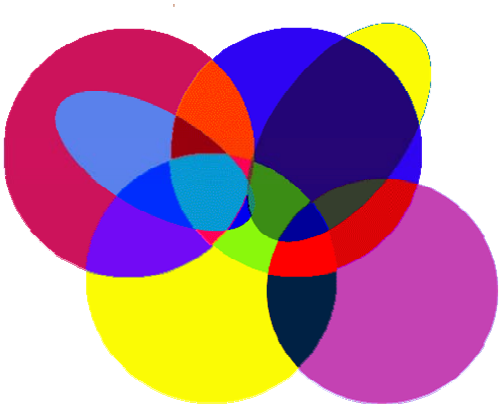

b) ***OSI Layer Isolation*** - Notional views that the layers of the OSI model are fully isolated from each other are obsolete. Net-centric applications and services require the ability to smoothly communicate with systems such as the routing infrastructure to control bandwidth allocation. Applications will also need to communicate with each other directly or through a mediation service to control priority and bandwidth assignment. Deployment of a faster, better network with IPv6 will not comprehensively address the needs of network-centricity without significant work being completed up front to model the system behavior, or else, the networks will be improperly sized/designed. The promise of these systems is to efficiently and intelligently use the vast capabilities of the network. However, new methods of communicating between the application and infrastructure layers, prioritization, scheduling and modeling are absolutely required in order to optimize the behavior of these complex systems.

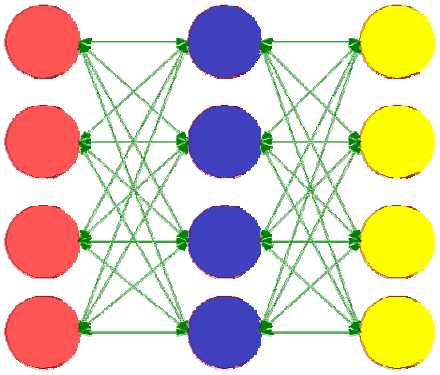

Problems with Application Layer Architecture – A subtle but significant problem exists when trying to efficiently distribute service architectures across complex networks. The principal behind the deployment of services architectures is to allow effective resource reuse between applications. In the net-centric model, functional capabilities traditionally viewed to be part of the application will be deployed as re-entrant services that are not owned by a specific application. This infers a new model for application support, one that substantially alters who is responsible for what portions of the networked application. Without network-based services that can be efficiently shared between applications and be easily reconfigured, network-centric systems will simply be too complex and expensive to deploy and maintain.

The question remains, however; who will be responsible for provisioning, maintaining and guaranteeing service levels for the shared services? Although these shared, distributed services lie well within the application layer, they are not owned by individual applications or users. The distribution and reuse of the services is more symmetric to the task of network infrastructure deployment than it is to "application" deployment. This requires a distributed, infrastructure-oriented support capability to implement core services. The roles of traditional network providers such as MILSATCOM must be redefined to include support for the services architecture when appropriate.

Intelligent Information Management: Distributed Content Caching, Data Object Routing and Distribution, Transaction Recovery, Fail-over Redirection and Persistence Management – The task of connecting applications and user communities in meaningful ways when applications are not necessarily aware of each other's existence is a difficult task. Information will be shared based on its characteristics and made available to systems and users based on their characteristics and missions. The challenge is to get the right information, to the right people, at the right time. The fabric of the network itself has to be intelligent to accomplish this goal and significant information handling capabilities must be built into the network fabric itself.

One of the difficulties with services-based architectures is that behavior of the services has to be predictable and symmetric to allow seamless interoperability and reliable performance. If the system behaves differently depending on where you are, or how you are connected, there will be problems with the architecture. Disparate, widely separated systems will not have a predictable, consistent or accurate view of what is happening across the network based on inaccurate assumptions about system status and information flow. One of the biggest complications is that these systems will have to efficiently handle highly asynchronous behavior with intermittency, attack disruption, network discontinuity and system interruptions creating an environment with ever-changing asymmetries. The only viable means of handling these problems is to create a network fabric that manages basic information movement tasks and guarantees certain performance characteristics so that edge systems can be reliably designed to interface across such a dynamic systems environment. Further, the concept of managing traffic at the object or information level heavily impacts the traffic that is required to move information across the network to multiple destinations.

If a point-to-multipoint distribution of information is required, it does not necessarily require that all information be transmitted directly form source to user. When one system requests an object, another system may need it in a few milliseconds or within the next few hours. To optimize network utilization the network services fabric must manage:

- delivery guarantee,
- persistence management,
- information caching,
- routing,
- recovery,
- restart,
- synchronization, and
- IA at the information level.

As a result, redundant transmissions can be eliminated and reliability enhanced. This requires a level of separation, open connection specification, and abstraction not commonly implemented in systems. The problem within the net-centric model is that the application

itself will become an ad-hoc combination of systems, services, users and data all needing common well-defined behavior across the network to allow the application designer to be able to implement applications that behave in a predictable fashion.

c) ***Prioritization, Scheduling, Application Mediation, Efficient Inter - Application Multiplexing and Coordination –*** One aspect of complex application management that must be controlled is the assertion of privilege. The right to operate now, the ability to reserve capacity, the right to schedule future activity and the need for near real-time feedback concerning the status of an application's request for bandwidth, must all be accommodated. This is not as simple as writing QoS bits into an IPv6 packet. The basic shared architecture of the application environment must allow the coordination of applications at a granular level. There must be a shared process for certifying the identity of the requesting system and an ability to differentiate requests coming from a particular system based on the data flow's particular requirements. Support for granularity of permission levels within an application domain will allow the application to step down the priority for certain functions. Packet level assertion will make the network susceptible to denial of service attack from systems that spoof address and QoS information. In this complex net-centric world, information and application level IA constructs must be used to allow the network to function properly. Otherwise, the network would require the continued arbitrary isolation of bandwidth which would disallow the system from taking advantage of even a reasonable percentage of the bandwidth that could be released from fixed, but unused, virtual pipes at any given time.

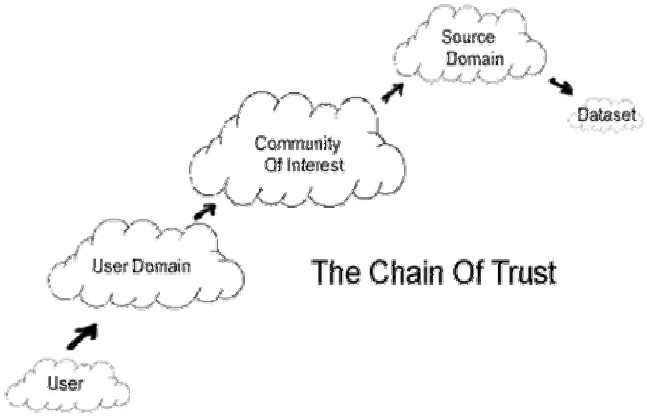

Cross-Domain and Object-Based Information Assurance – The new architecture requires that information itself be protected. The intersection of vast communities distributing mixed content across networks with variable security models requires that the IA design protect content at the information level. Such systems will have to support transport of information across boundaries without risk of improper access or disclosure of information. Multiple-level security will have to be applied to content so appropriate extracts or versions of information are disclosed only to the appropriate users or systems. Content protection is not the only aspect of this that will be critical, new architectures will be required to maintain:

- chain of custody and trust,
- instruction authentication,
- permission mapping, and
- nested/compound signatures.

The IA solutions will have to be able to map users, application domains, communities of interest, source domains and data systems by dynamically controlling access and system permissions. This model will require a level of granularity not previously contemplated. The basic IA architecture will have to operate with variable levels of assuredness based on operational and security tradeoffs. Some aspects of the IA model will require distributed services, others will require new hardware that can adequately guarantee the IA functions. In these environments it will be impossible to gain the desired operational capabilities if different systems and users with varying authorities are not allowed to tie together.

Identity on the Move and Distributed Identity Management – The participation in these networked systems is likely to be ad-hoc and situationally dependent. Participants are going to be on the move, connecting and then dropping connections in one location, and then reasserting connections in another location in a completely different context. The environment itself will have to deal with identifying participants, network context, location, equipment capabilities and a number of other parameters to properly enable the network applications. To this end, standards for asserting this information to the networked applications will have to be developed and broadly implemented. The further requirement for uniformity across

systems will drive toward inter-system and inter-application standards for these processes.

Hazard of Blending Layers - Use of one layer of the OSI model to implement the functionality of another will inevitably lead to inefficient network design. An example of this emergent difficulty is the work in progress regarding the use of IPv6 as the control point for Quality of Service (QoS), Class of Service (CoS) and RSVP control, without establishing how programs and systems should interface with these transport attributes within the network. Such assertion is by nature an application specific function, but no comprehensive design has been accepted as to how an application is to request a particular QoS or CoS level, or how the assertion of such a request is to be secured and authenticated. Uncontrolled assertion of such privileges would make the network susceptible to denial of service attack at the transport layer through malicious manipulation of the QoS and CoS requests. Just as important, it will be vital for the applications to be able to negotiate the proper access authority balanced against the current requirements of other applications and missions. Static assignment of such permissions would ignore the nature of warfare with constantly changing mission priorities as strategic and tactical situations change.

Object Transmission Modeling vs. IP Capacity - The traditional packet transport capacity and routing models for managing capacity utilization, network operations, inter-application coordination, and data transport are not sufficient to manage or properly design the network segments for network-centric applications. Further, technology for optimizing information flow through such networks must be deployed across the networks, not just at the edges. The choice of technical approaches to such optimization has significant hardware and software implications and substantially affect the security / information assurance architecture of the applications and the network. The choices will directly impact:

- the latency of information transmission,
- efficiency of data object dissemination,
- design of information assurance models to support a distributed role-based access control environment,
- requirements for granularity of data object attribution,
- support for full data multiplexing versus piping, and
- system capacity and security.

A simple example might make this concept more tractable. Take a network that has an effective IP packet transmission level of one million bytes per second. A network analysis would show that this would be plenty of capacity to transport data objects that are quite large. However, let's now say that the network suffers interruption every five to ten minutes because of tactical deployment problems and all transmissions of large files are made in a

session-controlled environment that uses simple FTP. Every time the network goes down, the sessions are lost and the transmission has to restart. Suddenly the effective capacity of the network to transport large data files is cut dramatically. Further, session-controlled streaming data cannot be transmitted across such a network for periods longer than the interruption frequency. Ad-hoc use of the network with session control and unscheduled access complicates this picture considerably. Use of synchronous session-oriented technology such as HAIPE will make the system extremely susceptible to transmission failure. Further, if the application layer penetrates multiple networks, it is the compounded reliability of all the network components in aggregate that determine the ability to get a transmission through the network. The important conclusion is that modeling network behavior requires much more than analysis of the network routing capabilities.

One needs to understand the impact various layers above the transport layer will have on the characteristic capacity of the network itself. Examples of where such application layer modeling would be important are:

a. Use of session-based Information Assurance,
b. Multi-hop network design considering the impact of timeout and message delivery guarantee on application architecture,
c. Intelligent edge-caching of certain types of content, and
d. Object persistence with network managed object recovery, reuse or rerouting.

The Alert Framework

Vision: Build an integrated command and control bridge, leveraging existing systems and new technologies to enable the transfer of actionable indications and warning information between DHS, DOD, AT/FP and public safety practitioners for the purpose of improving interdiction, containment and response in support of the national strategy for Homeland Defense/Security.

Mission: To provide a means of linking the multitude of communities and disparate Legacy systems that have Homeland Defense/Security functionality. Specifically assuring that the right information is delivered to right people at the right times, creating an effective means of sharing information without overloading the participants with irrelevant or untrusted data.

Capabilities (The system must):

Rapidly connect to disparate data sources while supporting changing connection specifications.

Support concepts of community ownership and data filtering to ensure data is released only as permitted, and has established trust and validity.
Maintain community mission integrity by active filtering of information.
Be easily deployed with a minimum footprint, impact on attached systems, and amount of operational overhead. Accommodate communication with a wide variety of established sensor, intelligence, case management and C2 systems, such as, DMIS, JPEN, and HSDN.

The ACTD provides the above requirements through its unique suite of technological capabilities; delivering assured communications, visualization, collaboration, command and control, situational awareness, information sharing, intelligent event management and alerting capabilities. The system provides an integrated environment for preventive preparedness, and improved consequence management/crisis response. The ACTD improves trust, interoperability, effectiveness and confidence of communications between non-government, and Federal, state and local agencies, including homeland defense/security first responders.

Framework Architecture

The framework implements a distributed agent based composite application architecture implementing shared services that interact with a common object definition for transporting alerts. The system is built on top of a rapid SOA application development platform that allows the flexibility to rapidly accommodate changing requirements. Structurally the framework implements a layered architecture for interconnections similar in nature to the OSI model. It is designed to enable seamless interaction between systems through the common object model. Internal systems use standards such as SOAP, CAP protocol for XML based alerts, JMS Publish and Subscribe as one means to manage message persistence, Java Security, and others but external connections can be configured to use almost any protocol. The used of this approach allows the incorporation of a new SOA or open transaction based connection or data consumer within an hour or so. External systems requiring more complex interfacing usually can be accommodated within a few days.

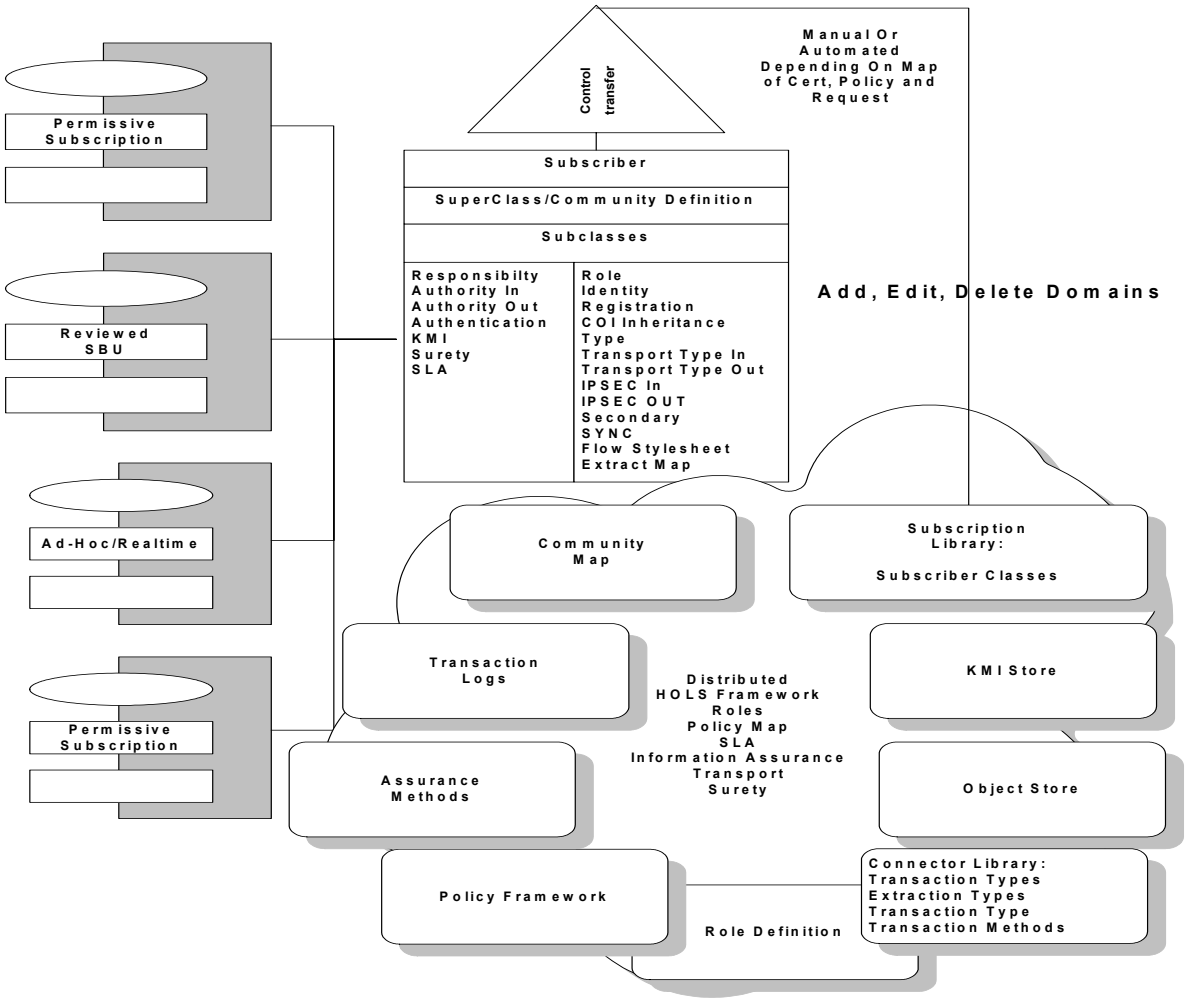

The Alert Framework is also designed to mix local and common policy control. This allows the local systems to control the flow of data in to and out of their systems. The common framework policy, on the other hand, insures only those parties with permission authority and valid credentials get access to the data to which they are supposed to have access. The implementation of the local interface layers are shown in the following figure. The open API allows the external system to plug into the framework at the point it can support. The goal was to build a systems that supported rapid connection building and rapid reconfiguration. The policy is implemented across the system as rule-sets that are not programming based. The rule-sets are exposed to the user community by a web application that gives access to participants by authority of those participants.

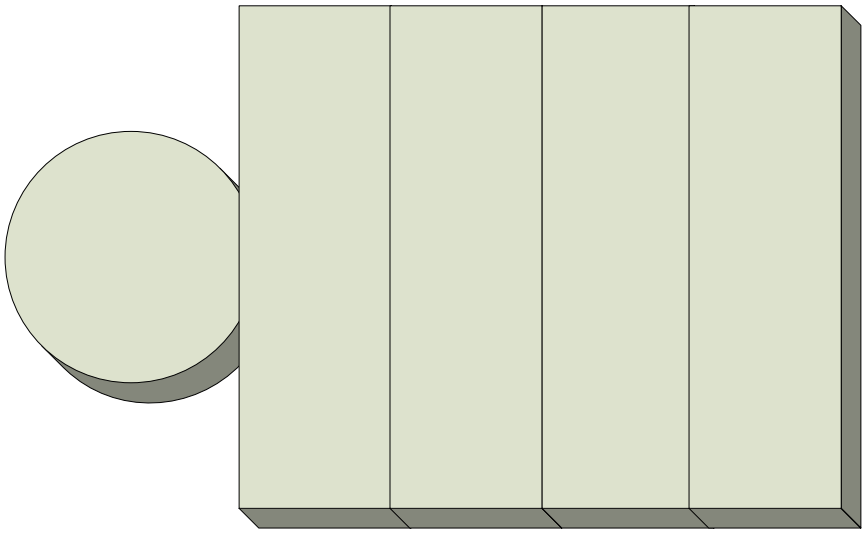

Within the primary framework common policy is enforced, workflow control is implemented and persistence management is provided. The workflow backbone allows near real-time reconfiguration of policy in the event of requirements changes. The system also supports centralized configuration of the framework with fully distributed agent based implementation. Workflow automation is also supported within the framework where specific state based manipulations of data objects are implemented by rules and/or proprietary code. By moving all transformations to the edge and using a common object model within the system architecture the framework is easily reused for different community alerting requirements,

Shared Community Framework

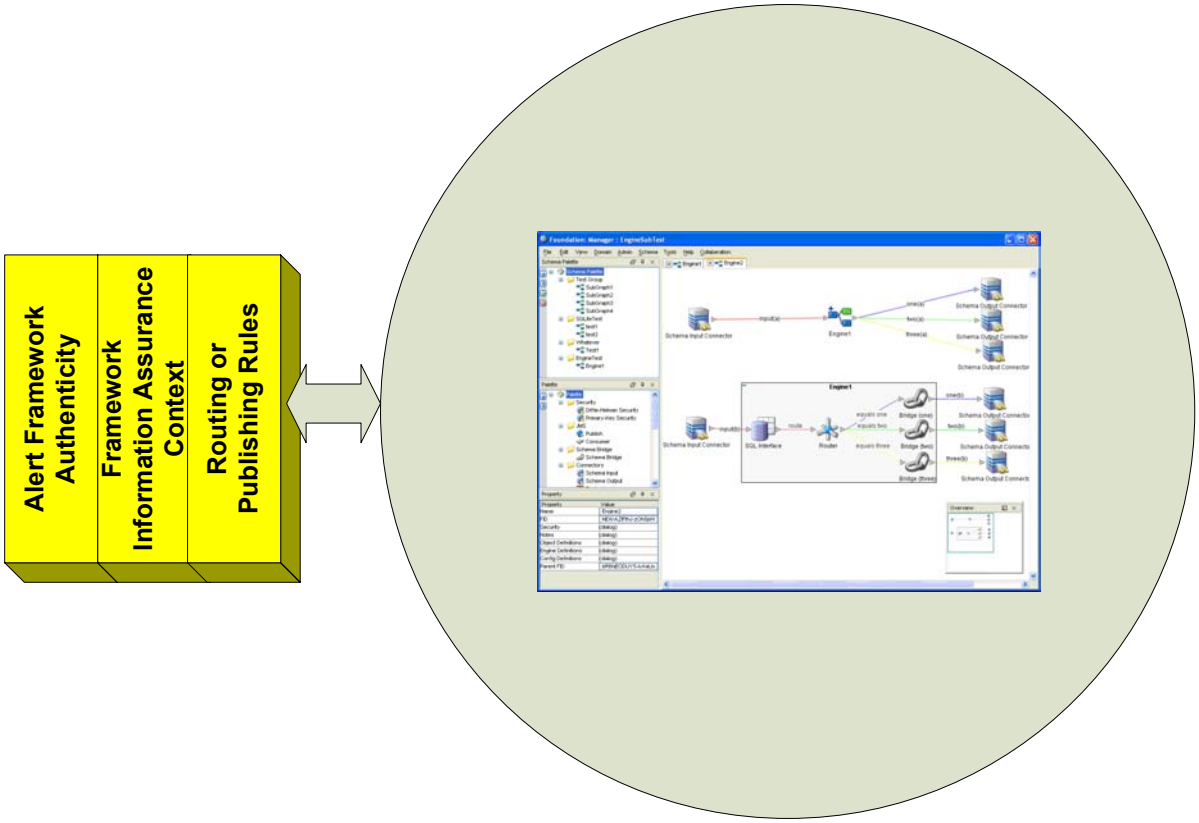

Next Steps

The Alerting Framework is a good first step in implementing effective, agile information sharing with extraordinary reach. The agent based environment scales very well but additional testing and simulation is required to examine large scale performance characteristics. It is also clear that there is not a single answer to all end user requirements and concepts of dynamic object transformation must be explored. Many of the capabilities that have been implemented go far beyond the current state of the art. It is important that the base level requirements are understood so that a unifying architectural approach to building solutions can be implemented. A properly layered applications architecture, combined with a broad understanding of network requirements, will be mandated to achieve the promise of the Alert Framework. Future work is slated for operational hardening, additional Information Assurance work and accreditation, and incorporation of additional communities.