

A Scalable Architecture for Improving the Timeliness and Relevance of Cyber Incident Notifications

James L. Miller, Robert F. Mills and Michael R. Grimaila
Department of Electrical & Computer Engineering
Air Force Institute of Technology
Wright-Patterson AFB, OH 45433-7765 USA
{James.Miller,Robert.Mills,Michael.Grimaila}@afit.edu

Michael W. Haas
711th Human Performance Wing
Air Force Research Laboratory
Wright-Patterson AFB OH 45433 USA
Michael.Haas@wpafb.af.mil

Abstract — The current mechanics of cyber incident notification within the United States Air Force rely on a broadcast “push” of incident information to the affected community of interest. This process is largely ineffective because when the notification arrives at each unit, someone has to make a decision as to who should be notified within their unit. Broadcasting the notification to all users creates noise for those who do not need the notification, increasing the likelihood of ignoring future notifications. Selectively sending notifications to specific people without *a priori* knowledge of who might be affected results in missing users who need to know. Neither of these approaches addresses the passing of notifications to downstream entities whose missions may be affected by the incident. In this paper, we propose a modular, scalable, cyber incident notification system concept that makes use of a “publish and subscribe” architecture to assure the timeliness and relevance of incident notification. Mission stakeholders subscribe to the status of mission critical information resources (external and internal) and publish their own mission capability allowing other units to maintain real-time awareness of their own dependencies. We contend that this architecture is a significant improvement over current methods by making direct connections between mission stakeholders and their dependencies and eliminating multiple levels of human processing, thereby reducing noise and ensuring relevant information gets to the right people.

Keywords- *CIMIA, cyber incident notification, situational awareness, mission assurance*

I. INTRODUCTION

Information is a critical asset in the operation and management of virtually all modern organizations [1][2][3]. Organizations embed information, communication, and networking technologies into their core mission processes as a means to increase their operational efficiency, exploit automation, reduce response times, improve decision quality, minimize costs, and/or maximize profit [4]. Military organizations are extremely dependent on access to information that is collected, processed, analyzed, distributed, and aggregated to support situational awareness, operational planning, intelligence collection and analysis, and command

decision making [5]. The increasing dependence upon information technology has resulted in an environment where an information incident (e.g., the loss or degradation of the confidentiality, availability, integrity, non-repudiation, and/or authenticity of an information resource or flow) can result in significant mission degradation or failure [6][7][8][9][10]. Research shows that most U.S. Federal organizations fail to implement key elements of established security guidelines such as providing a framework and implementing a continuing cycle of information security management activities including assessing and managing risk [11]. Even when an organization develops and maintains a robust security capability, it is inevitable that the organization will experience an information incident. When this occurs, it is important to notify the decision makers within organizations whose mission functions, processes, and/or tasks are critically dependent upon the affected information in a timely manner so they can take appropriate contingency measures [12].

Cyber Incident Mission Impact Assessment (CIMIA) is a research program whose stated goal is to provide decision makers at all levels of the military enterprise with timely and relevant notification of cyber incidents and expedited access to an assessment of how that incident impacts their mission or missions [13]. CIMIA relies on multiple disciplines to include situational awareness, risk management, mission representation, mission impact estimation, and incident notification. This paper is proposes the development of a modular, scalable publish and subscribe architecture as a means to enhance mission assurance by providing automatic, timely, and relevant incident information to people and organizations who depend on cyber resources. While the focus of the paper is on enduring missions, the discussion and results can apply equally as well to discrete missions.

The remainder of this paper is organized as follows: in Section II, a working definition of missions and enduring missions is provided; in Section III, we discuss the shortcomings and inadequacy of the existing incident push notifications; in Section IV, we introduce targeted automatic pull notifications and present an architecture of how to get

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE APR 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011	
4. TITLE AND SUBTITLE A Scalable Architecture For Improving The Timeliness And Relevance Of Cyber Incident Notifications				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology, Department of Electrical & Computer Engineering, Wright-Patterson AFB, OH, 45433				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES IEEE Symposium on Computational Intelligence in Cyber Security (CICS), Paris, France, 11-15 April 2011					
14. ABSTRACT The current mechanics of cyber incident notification within the United States Air Force rely on a broadcast ?push? of incident information to the affected community of interest. This process is largely ineffective because when the notification arrives at each unit, someone has to make a decision as to who should be notified within their unit. Broadcasting the notification to all users creates noise for those who do not need the notification, increasing the likelihood of ignoring future notifications. Selectively sending notifications to specific people without a priori knowledge of who might be affected results in missing users who need to know. Neither of these approaches addresses the passing of notifications to downstream entities whose missions may be affected by the incident. In this paper, we propose a modular scalable, cyber incident notification system concept that makes use of a ?publish and subscribe? architecture to assure the timeliness and relevance of incident notification. Mission stakeholders subscribe to the status of mission critical information resources (external and internal) and publish their own mission capability allowing other units to maintain real-time awareness of their own dependencies. We contend that this architecture is a significant improvement over current methods by making direct connections between mission stakeholders and their dependencies and eliminating multiple levels of human processing, thereby reducing noise and ensuring relevant information gets to the right people.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

those notifications to enduring mission owners; in Section V, we provide a brief scenario of how such a system might relay information to the spider web like structure that exists in enduring missions; in Section VI, we discuss related works; in Section VII we discuss limitations with the proposed approach; in Section VIII we present conclusions and discuss future work.

II. MISSIONS

Department of Defense Joint Publication 1-02 defines a mission as, “The task, together with the purpose, that clearly indicates the action to be taken and the reason therefore,” and, “In common usage, especially when applied to lower military units, a duty assigned to an individual or unit; a task.” Smaller missions combine to form the basis of a larger mission from the smallest one person work center at a remote detachment to the Pentagon. This definition tends to refer to discrete tasks which have a definite start and stop criteria [14].

The concept of an enduring mission is not directly defined in Joint Publication 1-02, Department of Defense Dictionary of Military and Associated Terms. However, the term “enduring” is used within many other definitions within the publication. As an example, in the definition of a planning team it states, “The planning team is not enduring and dissolves upon completion of the assigned task” [14]. For the purpose of this paper, an enduring mission is one in that continues without a definite end point.

Most support missions and a significant percentage of operational missions on a particular installation (e.g. an army post) are enduring missions. Whether it be operating and maintaining the installation’s communications infrastructure or providing meals at a dining facility, enduring missions continue indefinitely. There are course corrections as consumer needs and priorities change, but the primary *raison d’être* for the mission continues for an extended period of time. These enduring missions typically support other enduring missions and may also directly or indirectly support discrete missions such as a convoy mission or an airlift mission.

Not all enduring missions are directly hierarchical in nature. Military organizations typically have “Additional Duties” which cover areas such as safety, security, or readiness within a unit. In this case, a commander appoints two or more individuals and those individuals report directly to the commander on that duty instead of following their normal hierarchical chain of command. Often an additional duty may require additional and recurring training that is specific enough to the duty that the personnel who work with this individual in their primary duty may be unable and/or prohibited from performing the additional duty should the appointed person be absent. Those individuals given an additional duty may or may not work together in their primary enduring mission. As implied by the name, the individuals so appointed still have a primary duty to which they are assigned, but this additional duty is done concurrently with their primary duty. With increased reliance on cyber assets in all missions, it is likely that those assigned these additional duties will access cyber assets for both their primary and additional duties.

Commanders at all echelons in all branches of the military are responsible for “everything their command does or fails to

do” [15]. Commanders delegate their authority (though not their responsibility) both hierarchically and non-hierarchically as appropriate to the need and in accordance with governing regulations or instructions. That delegation of authority also requires the communication of commander’s intent—which per Joint Publication 1-02 means “...concise expression of the purpose of the operation and the desired end state...” [14]. Whenever the commander’s authority is further delegated, there is a responsibility of the delegator to communicate the commander’s intent as it relates to that particular delegation of authority.

For purposes of this paper, an enduring mission can fall within or across three categories:

- A hierarchical unit performing a function or functions to which a person or persons have been delegated authority and provided with commander’s intent.
- Non-hierarchical (additional) duties in which authority of the mission has been delegated and commander’s intent has been communicated but is done outside of the normal hierarchical structure (i.e. chain of command).
- Operation and maintenance of cyber resources, such as circuits, databases, or servers, which often fall somewhere between hierarchical and non-hierarchical duties. Maintaining these items often includes multiple entities (i.e., the actions of personnel are required at both ends of a long-haul circuit to maintain its operation as well as the actions of personnel along the path of the circuit).

Leaders who possess that delegation of authority make decisions based on their understanding of commander’s intent and on their situational awareness. For purposes of this paper, Situational Awareness (SA) will follow Endsley’s SA model which includes: “... the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future.” Mission stakeholders operate between Levels 2 and 3 in Endsley’s model [16].

III. PUSH NOTIFICATION PROCESS

The default standard of providing Level 1 SA regarding cyber incident to users of information systems has been through a push notification. When an incident occurs, the entity that discovers the problem sends out a message to those who they think are or will be affected by the incident.

A working example of this push method is used by the United States Air Force. Cyber incident notifications are governed by Air Force Instruction (AFI) 33-138, Enterprise Network Operations Notification and Tracking. Notification of cyber incidents are done through Command, Control, Communications, and Computer Notices to Airmen (C4-NOTAM’s) and Time Compliance Network Orders (TCNO’s). A C4-NOTAM is defined in the instruction as “...closely related to TCNOs with the primary difference being that they are informative in nature and are not used to direct actions. They are used by all organizations within the AFNETOPS (Air Force Network Operations) hierarchy. They

are the primary means of disseminating network information that does not require specific actions to be taken, or compliance to be tracked...” By inference, TCNO’s are C4-NOTAM’s which require action compliance accountability [17].

Previous research identified five limitations in the existing notification processes used by the United States Air Force [13].

- Data providers are not able to identify critically dependent downstream consumers
- Failures exist in internal organization communication
- A functional automatic notification system does not exist
- Organizations and their missions are dynamic
- Criticality determination is organization dependent

Grimaila et al. [13] recommend automatic notifications as one of three potential ways to improve on the current status quo. The ability for units to automate incident notifications when resources or capabilities they depend upon are affected, provides valuable situational awareness that supports mission assurance.

C4-NOTAMs are text messages sent from one echelon to another within the AFNETOPS. They are labeled by type (Informative, Unscheduled Event, Scheduled Event, and Summary) and are distributed via a means consistent with their sensitivity. Upon receipt at each echelon, regardless of whether the message is traveling up or down the echelons, a human has to decide what to do next [17]. As long as this message does not require distribution outside of the network operations community, the flow of communications is reasonably straight forward, but information often flows at the speed of human receipt and processing. Alberts & Hayes [18] would describe the C4-NOTAM process as a series of pushes, with each push introducing a delay and an opportunity for the information flow to stop altogether.

When the notification leaves the network operations community to any of the other user communities (operations or logistics as an example), problems start to occur. As noted by Hale et al. [19], there is a chasm that exists between the network operations communities and the communities they support. The network operations community is responsible for all aspects of C4 systems but operations, as an example, sees C4 capabilities as a utility. When a cyber incident occurs, network operations personnel are tasked with collecting mission impact data from the user communities for assessment at higher headquarters. When it reaches higher headquarters, a network operator at that echelon briefs the commander in charge, usually with representatives of the other user communities present for analysis and opinion.

Tinnel et al. [20] speak to this kind of functional gap and asymmetric dependency between operations and information technology personnel. The same personnel who are trying to defend an active attack are simultaneously attempting to perform battle damage assessment. And rather than just collecting the data, often it is network personnel who are also trying to measure that damage because the users of the affected

resource(s) do not understand how important the resource(s) are to their ability to accomplish the mission.

But even more troubling, when the information needs to leave the network operations community and get to the other communities, there is no certainty that the message is going to be received by those who are affected by the incident. Grimaila et al. [13] classify this issue as a failure of internal organizational communication. When the C4-NOTAM reaches a base level network control center (NCC), an individual sans *a priori* knowledge of the plethora of missions on a base has to decide who to send the notification to.

Sending that notification to all base network users becomes noise when only a few individuals need the notification, making it possible that future notifications will be ignored [21]. If instead that same person tries to narrow the scope to just what is viewed as the target audience for the notification may then fail to notify parties who need to receive the notification. This is referred to as unintended message filtering [13].

Another problem with relying on this notification scheme is there is no certainty that even if the right mission is picked and the right person is picked that the “right person” will be present to receive the message. The standard in-garrison work week for most Airmen hovers around 40 hours, yet many missions carry on at all hours of the day and night on weekdays, weekends, and holidays. Issues that happen during the “normal” duty day may still not be picked up on a perfectly selective pushed message due to the intended recipient being gone due to illness, deployment, meeting attendance, meal break, or any number of other situations

Grimaila et al. [22] discuss a hypothetical convoy operations scenario as an example of how push notifications do not work effectively. In this scenario a database server used for mission planning is compromised. This database contained information regarding convoy operations. Once it was realized that the confidentiality of the database could have been and likely was compromised, a notification process was started. But for many reasons, the notification took longer than it should have and ultimately did not reach mission stakeholders who most needed the information. In this scenario, the result was repeated ambushes of the convoy.

An alternative push method would be to put the onus of user notification in the hands of the cyber resource provider. This would require that the provider knows who all of the consumers of his or her resource is and has good point of contact information for multiple personnel at the mission that is supported by this resource. As was a lesson learned of at least one of the authors of this paper during the run up to January 1st, 2000, maintaining such a list in a quasi-static environment was next to impossible during a time when changes to standard operating procedures were nearly frozen until the so-called “danger dates” had passed. Add more than a decade and our networks are much larger and more dynamic, making this approach that much more difficult.

IV. TARGETED AUTOMATIC NOTIFICATIONS

Targeted automatic notification of cyber incidents is the ideal outcome. This would reduce the noise associated with notifications that are not relevant to the recipients. Pushing

automatic notifications is problematic based on an inability to maintain a relevant list of potential recipients. A better way would be to move from pushing notifications to publish and subscribe, or pulling notifications with an agent-based scheme.

We propose an architecture that provides the ability for mission stakeholders to publish the status of their mission while also having the ability pull the status of the resources that the mission depends on.

This mission capability status is published at the will and discretion of, and is defined by the mission stakeholders as it relates to their understanding of commander's intent. Whether the capability depended on from that asset is cyber-related or more general is of little consequence.

Figure 1 illustrates a notional 3 level architecture and the paths of communication between these levels.




Figure 1: A Notional Three Level Cyber Incident Notification Architecture

The architecture starts at a top level server. This top level server stores primarily the status of cyber assets that are used at more than one base or base-equivalent. It is located within an organization that has an adequate amount of oversight of current cyber threats and attacks, and personnel at that organization may manually update the status if there is a known or suspected cyber incident occurring. A top level server receives updates from local level servers when status changes occur.

Local level servers act as the middle man between the user level agents and the top level server. Local level servers reside on a local base's communications infrastructure and can be operated locally or remotely. The server is responsible for:

- Storing local cyber asset status
- Storing local enduring mission status
- Storing local enduring mission configuration files
- Performing periodic queries to local cyber assets (as registered with the server) to ensure availability (similar to Command and Control Remote Monitoring System (C2RMS) [23])

- Performing manual status updates/overrides of local cyber assets when warranted
- Requesting and storing remote cyber asset status from the top level server
- Publishing the status of local cyber assets to the top level server

The local server requests a new status from the top level server when it receives a request from a user and that user needs a more recent status check than the base level server has on record. When the new record comes back from the top level server, the base-level server then communicates the most recently retrieved status.

In this notional architecture, there is adequate flexibility to allow for additional levels to exist between the top level and the local level servers. There are situations tied to geography, distance, and density of use where an intermediate level would be beneficial, similar to the scheme employed for domain name service (DNS) Servers [24]. The intermediate servers would contain a subset of what is cached at the top level server and would pass data between the top level and local level servers. The actual mechanics is a topic left for future research studying the trade-offs between span of control and network overhead.

The user-level agent acts as a status checker for the user running the agent. This agent may monitor and update multiple enduring missions depending on the responsibilities of the user running the agent. For each enduring mission that user is responsible for, status reports are retrieved from the base-level server for that enduring mission's dependent missions. Those statuses are retrieved in a time interval commensurate with the importance of that dependent mission on the execution of the monitored enduring mission as identified by the dependent mission leaders based on commander's intent.

The time interval for performing checks may be overridden in instances where a new update is needed immediately but doing so should be discouraged through rule or mechanics. First, the potential exists for an intentional or unintentional denial of service attack by repeated out of band status checks. And secondly, the strength of periodic checks is that eventually these information pulls blend into the normal chatter of network traffic. An intruder monitoring network traffic should see no abnormal rise in traffic with strictly automatic notifications as opposed to the significant increase in traffic that would result from push notifications [13]. Additional requests may rise above the noise level to provide an intruder some insight that an issue exists.

The user-level agent generates an alert when there is a difference between the last two status messages received. The status message may reflect anything on a wide spectrum of events from a mission capability change or significant activity (SIGACT) report as defined by Joint Publication 1-02 [14] to something innocuous (limited personnel available for a period of time due to a military function). In the initial implementation, it is up to one of the mission stakeholders cognizant of their commander's intent to determine whether or not to update their status so that dependent missions can become aware of the change. This is an area where additional work could be done to interpret the status messages as they are

received and use case-based reasoning to quickly determine how the change in status has affected the mission before and provide the decision maker with a list of options.

The user-level agent also is the means for creating and modifying configuration files for each enduring mission. During initial setup, a leader can identify what it is that they use and set an alert commensurate with the importance to the mission. The mission stakeholder gets the identification information for their dependencies from the owners of that dependency. How public this information is remains the discretion of the providing mission stakeholder. As the mission evolves and the available tools change, modifications can be made to the configuration files.

Figures 2, 3, and 4 illustrate the importance of the configuration files as they relate to the user agent. Figure 2 shows the local level server is storing the configuration files for Individuals A, B, and C, all who have authority for Mission A. Their user agents all have Mission A loaded to them. Individuals A and B have been assigned to Additional Duties B and C respectively for which individually they have authority but their co-workers do not. Each individual has a slightly different view within in their user agent based on the missions they have authority for.




Figure 2: A View of a Notional Local Level Server's Stored Missions and Three Notional Users

In Figure 3, if an event caused an alert for the Mission A, all three individuals would see that alert. The alert would be triggered in the agent at the next update interval for each agent. Whether all three user agents would get the update at precisely the same time or within a window of time equal to the defined interval between checks is a detail left for implementation.

In Figure 4, if there was an event that caused an alert for the Additional Duty B, only individual A would get that alert. If individuals B or C had gotten the alert regarding Additional Duty B, that alert would have been noise as they have no authority over it.




Figure 3: The Local Level Server from Figure 2 With An Alert for Mission A




Figure 4: The Local Level Server from Figure 2 with an Alert for Additional Duty B

Subscription to resources is made through a process that starts with introspection on what is required for mission accomplishment. We take the position that successful leaders can identify most, if not all, of what they depend on to meet commander's intent. The method of finding the unique identifier for those depended-upon resources will vary based on the nature of the resource. Obtaining an identifier may be as easy as viewing a public website or as difficult as submitting a paper-based application to the mission owner. The method of obtaining the identifier would be based on the sensitivity or relative value of the resource in the eyes of the provider. Helper applications discussed later may be used to help identify overlooked mission dependencies.

V. ANATOMY OF A CYBER INCIDENT NOTIFICATION

In a hypothetical military organization at Main Operating Base Alpha (MOB-A), a power supply fails in a server operating a single database. The database serves a small subset of users at a number of military bases, to include an organization at Main Operating Base Bravo (MOB-B). The database has been identified as requiring incident notifications in accordance with applicable instructions. The owner of the

database recognizes the problem and the need to make the report.

In a traditional push communication method, the database owner notifies the local communications entity, the MOB-A Base Communication Center (BCC). The BCC pushes the message up to the next higher level of communications responsibility. The message is pushed up the chain twice more until it reaches the top level of communications responsibility (TLCR).

The TLCR is aware that the database could be used by other organizations but they do not know exactly who uses it. So the TLCR pushes a message to all of their immediate subordinate communications organizations. This push continues utilizing one push at a time down to all of the BCC's.

Provided that the MOB-B BCC gets the notification, a human there has to decide who to send the message to. As this is a reasonably specialized system, a message sent to all MOB-B users will be noise to all except the one or two that use the database. If the human instead decides to be selective as to whom they send it to, whether or not the dependent mission gets the notification will be purely on how much *a priori* knowledge that person has or how well they guess. If the guess was accurate and the message is received by the mission stakeholder, they can then assess how it will affect their mission. It is up to the mission stakeholder to then alert their downstream users with similar *a priori* knowledge of who might be affected by the resultant change in mission capability.

If instead our proposed architecture is used, the database owner uses their agent and publishes the status change. The local level server recognizes that it is a cyber asset and publishes the status change to the top level server.

Minutes later at MOB-B, a user agent requests the status of the cyber asset as it does periodically throughout the day. The local level server at MOB-B sees that the information it has stored is too old for the user need and requests a status update from the top level. The top level server responds with the most recent status which the local server echoes to the requestor and stores. Because this was a change in status, the user agent displays an alert.

With receipt and observation of the alert (Endsley's Level 1 SA), the stakeholders responsible for the enduring mission that requested the update need to decide what this means to them (Level 2 SA) as well as make any changes to how they are performing their mission to accommodate for the lost resource (Level 3 SA) [16]. If it will negatively affect their ability to perform their mission, they publish that update to the local server and that information propagates to downstream missions dependent on that particular enduring mission when their user agents request an update. Because the mission that used the database at MOB-B was not related to a cyber asset, the update would go to the local level server but not the top level server.

Alternately, if this had been a compromise of data integrity or confidentiality, an appropriate warning would go out as well. This could warn downstream users that there could be problems with previously provided products and appropriate actions need to be taken. In the aforementioned scenario of a convoy database confidentiality breach, an alerting system such

as this could save lives. Assisting the users with determining if/how the compromised information was used is a separate problem beyond the scope of this research.

VI. RELATED WORK

The concept of a central repository for distributing automatic notifications is not a new one. Stanley et al. [25] proposed a Mission Service Automation Architecture with a centralized database to track configuration of the network from which users with different levels of authority and responsibility of the network enterprise could share updates. Fenz et al. [26] proposes the concept of a central entity to collect software vulnerability notifications and distribute it to subscribers.

Publish and subscribe is also not a novel concept and is discussed at length in [18][27][28]. While this work does easily integrate the concept of smart pull [28], work could be done to do so. Publish, subscribe, and query is part of the Joint Battlespace Infosphere [29].

Our research is inspired by the work documented in C2RMS [23]. C2RMS within a Combined Air Operations Center (CAOC) performs pings and data inquiries to systems that the CAOC needs to perform their missions. When one of these checks fails, a previously generated alert appears telling of the failure and what it means to the mission. In making a data system available to respond to pings or data inquiries, it effectively is passively publishing its availability status and the act of pinging is a pull action. However, those notifications stop with the end user who was monitoring the resource. Future notifications occur outside of the C2RMS framework, and for a CAOC this works because all of the players and resources are more or less in the same room.

Much of the current research in this area of interest is focused on observing information flows in the network and mapping these flows to the missions and people consuming them. Camus [30] can perform this through comparing logs to Lightweight Directory Access Protocol (LDAP) inquiries. Mission Aware Reporting of Information Assurance for Airborne and Enclave Networks (MARIAAN) [31] attempts to bring mission events into the comment event expression format and then try to analyze these events to increase mission awareness. Various schemes such as information asset valuation [10] have been suggested as a way to tie network traffic with the missions they support. And while this is undoubtedly the future and goal to be strived for as a piece of the solution, it does not address the downstream users. And all have relied on push rather than pull.

Additionally, there are three major hurdles in the way to tie data with mission execution. The dynamic nature of both military missions and the personnel who have the authority to execute them creates a roadblock to finding the right person or persons at the right time. The environment of "do more with less" that has punctuated the United States Air Force's over 40% reduction in personnel since the end of the Cold War [32] or "do more without more" vision which has now entered our lexicon [33] has increased the reliance of self-service for tasks previously performed by administrative or contracted personnel, introducing non-mission noise onto network devices performing mission-related tasks on the Non-classified Internet

Protocol Router Network (NIPRNet). And in a NIPRNet environment where a significant portion of the network traffic is authorized morale-based traffic (i.e. Facebook), sifting through the noise to find all the key data runs the risk of inadvertent filtering.

Where this research is unique and novel is that our focus and concern is connecting the spider web of enduring mission owners with those that depend on them and a reliance on active publishing of status updates. We acknowledge and embrace the fact that enduring missions act as both consumer and supplier to additional enduring missions and our focus with this research is providing the communication pathways.

We also acknowledge achieving success with this architecture requires introspection by mission stakeholders to determine what their missions use and what is important. This is similar to knowing if the proper tools and technical orders available to fix an aircraft are available or if there are enough people to fully staff the dining facility for the mid-day meal. The primary difference is that this requires introspection of tools contained almost completely in the information domain rather than in the physical domain [27].

VII. LIMITATIONS

This research relies on the proposition that enduring mission leaders are aware of what systems and other dependent missions they need to execute commander's intent and how important those resources are to that execution. Not understanding what is required to perform a mission is not a technology problem—it is a people and process problem. Research such as that done by Milcord with Commander's Learning Agent could help populate the user agent to ensure that cyber assets are not forgotten [34].

We acknowledge that achieving Level 1 SA is only a step to achieving mission success. A timely warning light on a vehicle's dashboard is relevant in that it pertains to that particular vehicle, but it means little to someone who does not comprehend its importance. The aim of this architecture is to enhance timeliness and relevance by speeding the notification process and getting notifications to those who have requested them. It does not nor is intended to enhance the relevance of the message contents. Other research taking place at AFIT concurrently with this research is working towards validating that targeted notifications will help mission effectiveness and that case-based reasoning can further enhance the relevance of these warnings based on prior incidents.

We acknowledge that any increase in relevance of notifications is limited to a reduction in noise received as compared to the current e-mail method. Further research is required to compare notifications received with the actions taken in similar previous incidents to further improve relevance.

This architecture generally relies on an outside source to detect when there is a cyber incident. Limited capability exists to check for availability via a C2RMS-like mechanism. Otherwise the architecture acts only as a means of transporting notifications.

Full comprehension of the importance of an asset to mission impact is far less than precise, especially with the dynamic nature of the military. As has been said by many in many circumstances, no plan ever survives first contact with reality. That is where leadership and experience takes over and lessons learned are used to modify previous assumptions.

VIII. CONCLUSION

Cyber incident notification procedures, as currently employed as pushes, are slow and either unfocused or inaccurate. Mass notifications to unaffected or disinterested personnel creates noise and create the potential for unintended information filtering. Manual targeted notifications require a level of *a priori* knowledge about usage and criticality that does not exist as it is difficult to obtain and maintain. Enduring missions would benefit from the development of a publish and subscribe architecture. This would enable mission stakeholders to identify what it is that is important to them and remove the middle layers of notification that exist in a push architecture. Speeding the notification process could enable quicker mission impact assessment by mission stakeholders, provide a conduit to notify downstream dependent missions, and free information technology personnel to fight through attacks and restore systems to full operational capability.

The architecture presented in this paper is a concept for a rudimentary decision support system. We believe that such a system will provide better Level 1 SA [16] than is currently available using the existing push incident notification process. This work is part of the ongoing CIMIA research program sponsored by the Air Force Research Laboratory.

Work is still needed on many fronts with this notification architecture. Format, content, and data structure of the status messages is a major detail that needs to be worked out. Applying Occam's razor to the problem would suggest implementing small text files with a minimum of information and a DNS-like serving mechanism would be quick and unobtrusive on the network. One of the next milestones is building a simulation and/or a prototype to demonstrate potential functionality.

Keeping this much information in one centralized location may also run afoul of Operations Security (OPSEC) practices. OPSEC may dictate that this system be kept on a higher security classification network. Provided that these details could be worked out, then there is also future work that can be done to utilize this information in a more automatic way. A related research activity is underway at Air Force Institute of Technology to apply case based reasoning to enhance the relevance of the incident notification process to improve mission assurance by informing commander's decision making process.

The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

IX. REFERENCES

- [1] Denning, D. (1999) "Information Warfare and Security," Upper Saddle River, NJ, Pearson.
- [2] Pipkin, D.L. (2000) "Information Security Protecting the Global Enterprise," Hewlett-Packard Company.
- [3] Fortson, L.W. and Grimaila, M.R. (2007) "Development of a Defensive Cyber Damage Assessment Framework," Proc. of the 2007 International Conference on Information Warfare and Security (ICIW 2007); Naval Postgraduate School, Monterey, CA.
- [4] Grimaila, M.R. and Fortson, L.W. (2007) "Towards an Information Asset-Based Defensive Cyber Damage Assessment Process," Proc. of the 2007 IEEE Computational Intelligence for Security and Defense Applications (CISDA 2007); Honolulu, HI, pp. 206-212.
- [5] JP 3-13 (2006) "Joint Publication 3-13: Information Operations," Joint Chiefs of Staff, United States Department of Defense, 13 February 2006, http://www.dtic.mil/doctrine/jel/new_pubs/jp3_13.pdf.
- [6] Jajodia, S., Ammann, P., and McCollum, C.D. (1999) "Surviving Information Warfare Attacks," IEEE Computer, Vol. 32, No. 4, pp. 57-63, April 1999.
- [7] Finne, T. (2000) "Information systems risk management: Key concepts and business processes," Computers & Security, 19, 3, 234-242, 2000.
- [8] Anderson, E., Choobineh, J., and Grimaila, M.R. (2005) "An Enterprise Level Security Requirements Specification Model," Proc. of the 38th Annual Hawaii International Conference (HICSS 2005), Jan. 2005, pp. 186-196.
- [9] Sorrels, D., Grimaila, M.R., Fortson, L.W., and Mills, R.F. (2008) "An Architecture for Cyber Incident Mission Impact Assessment (CIMIA)," Proceedings of the 2008 International Conference on Information Warfare and Security (ICIW 2008), Peter Kiewit Institute, University of Nebraska Omaha.
- [10] Hellesen, D., Grimaila, M.R., Fortson, L.W., and Mills, R.F. (2008) "Information Asset Value Quantification," Proc. of the 2008 International Conference on Information Warfare and Security (ICIW 2008), Peter Kiewit Institute, University of Nebraska Omaha, April 24-25, 2008.
- [11] Wilshusen, G.C. (2009) "CYBERSECURITY: Continued Efforts Are Needed to Protect Information Systems from Evolving Threats," United States Government Accountability Office, <http://www.gao.gov/new.items/d10230t.pdf>
- [12] Grimaila, M.R., Fortson, L.W., and Sutton, J.L. (2009) "Design Considerations for a Cyber Incident Mission Impact Assessment (CIMIA) Process," Proc. of the 2009 International Conference on Security and Management (SAM09), Las Vegas, Nevada, 2009.
- [13] Grimaila, M.R., Schechtman, G., and Mills, R.F. (2009) "Improving Cyber Incident Notification in Military Operations," Proc. of the 2009 Institute of Industrial Engineers Annual Conference (IERC 2009), Miami, FL.
- [14] JP 1-02, "Joint Publication 1-02: Department of Defense Dictionary of Military and Associated Terms," Joint Chiefs of Staff, United States Department of Defense, 12 April 2001 (As amended through 19 August 2009), http://www.dtic.mil/doctrine/jel/new_pubs/jp1_02.pdf.
- [15] AR 600-20 (2010) "Army Regulation 600-20: Army Command Policy," Headquarters, Department of the Army, http://www.army.mil/usapa/epubs/pdf/r600_20.pdf.
- [16] Endsley, M.R. (1995) "Toward a Theory of Situation Awareness in Dynamic Systems," Human Factors Journal, 37(1), pp. 32-64.
- [17] AFI 33-138 (2005) "Air Force Instruction 33-138: Enterprise Network Operations Notification and Tracking," Headquarters Department of the Air Force, <http://www.e-publishing.af.mil/shared/media/epubs/AFI33-138.pdf>.
- [18] Alberts, D.S. and Hayes, R.E. (2003) "Power to the Edge: Command... Control... in the Information Age," US DOD Command and Control Research Center Publications.
- [19] Hale, B., Grimaila, M., Mills, R., Haas, M., and Maynard, P. (2009) "Communicating Potential Mission Impact Using Shard Mission Representations," Proceedings of the 5th International Conference on Information Warfare and Security, Wright-Patterson AFB, Ohio
- [20] Tinnel, L.S., Saydjari, O.S., and Haines, J.W. (2003) "An Integrated Cyber Panel System", Proceedings of the 2003 DARPA Information Survivability Conference and Exposition, Vol. 2, pp. 32-34.
- [21] Libicki, M.C. (2007) "Conquest in Cyberspace: National Security and Information Warfare," RAND Corporation
- [22] Grimaila, M.R. (2008) "Improving the Cyber Incident Mission Impact Assessment Process," Proc. of the Cyber Security and Information Intelligence Research Workshop (CSIIW 2008), Oak Ridge National Laboratory, Oak Ridge, TN.
- [23] Jos, B. and Culbertson, T. (2006) "Leveraging Net-Centric Monitoring Techniques with Information Fusion to Increase US Air Force Information Dominance," Military Communications Conference (MILCOM 2006), Washington, DC, pp. 1-6.
- [24] Kurose, J.F., Ross, K.W., (2010) "Computer Networking: A Top-Down Approach (Fifth Edition)," Addison-Wesley.
- [25] Stanley, J.E., Mills, R.F., Raines, R.A. and Baldwin, R.O. (2005), "Correlating Network Services with Operational Mission Impact," Military Communications Conference (MILCOM 2005), Atlantic City, New Jersey, Oct 2005, pp. 1-7.
- [26] Fenz, S., Ekelhart, A. and Weippl, E. (2008) "Fortification of IT security by automatic security advisory processing", Proceedings of the 22nd International Conference on Advanced Information Networking and Applications (AINA2008), IEEE Computer Society, Los Alamitos, CA, USA, pp. 575-582.
- [27] Alberts, D.S., Garstka, J.J., Hayes, R.E., and Signori, D.T. (2001) "Understanding Information Age Warfare" US DOD Command and Control Research Center Publications
- [28] Alberts, D.S., Hayes, R.E. (2006) "Understanding Command and Control" US DOD Command and Control Research Center Publications
- [29] Combs, V., Hillman, R., Muccio, M., and McKeel, R. (2005) "Joint Battlespace Infosphere: Information Management within a C2 Enterprise" The Tenth International Command and Control Technology Symposium (ICCRTS), Virginia Beach, VA.
- [30] Goodall, J.R., D'Amico, A. and Kopylec, J.K. (2009) "Camus: Automatically Mapping Cyber Assets to Missions and Users," IEEE Military Communications Conference, Boston, MA.
- [31] Heinbockel, W., Kertzner, P., McQuaid, R. (2010) "Providing Mission Assurance for Airborne Networks", IEEE International Conference on Privacy, Security, Risk, and Trust, Minneapolis, MN
- [32] Thatcher, S. (2010) "The Secret of Doing More With Less", <http://www.mcconnell.af.mil/news/story.asp?id=123214401>, retrieved 20 September 2010.
- [33] Carter, A.B. (2010) "Better Buying Power: Mandate for Restoring Affordability and Productivity in Defense Spending," Department of Defense.
- [34] Milcord (2010) "Commander's Learning Agent", http://wiki.milcord.com/index.php/Commander's_Learning_Agent, retrieved on 20 September 2010.