

Technology solutions for coalition operations

Douglas Wiemer

Abstract— As computer technology has advanced, information processing in command, control, communications, computers, intelligence, surveillance and reconnaissance (C4ISR) systems has become highly complex. The information processed by these systems is usually of a very highly sensitive nature and is entered into specific systems that are physically isolated from each other. The physical isolation of these systems makes it cumbersome to exchange information between systems. The result is inefficient sharing of sensitive information in situations where timeliness of exchange could be a life or death reality. Since the mid 1990's, increasing efforts have been placed on improving coalition operations. Many systems have been created with the goal to improve the sharing of information and collaborative planning across coalition boundaries. The usability of these systems have had mixed levels of success and improvements will always be necessary. This paper will briefly describe three advances in telecommunications technology that could be leveraged to significantly improve coalition operations. These technologies are; the session border controller (SBC), advances in pattern matching technology, and multi-protocol label switching (MPLS).

Keywords— C4ISR, coalition, CCIS, command and control, information technology, military, defence, telecommunications, operations, SBC, session border controller, content inspection, MPLS, multi-protocol label switching, pattern matching, trusted downgrade, explicit route.

1. Introduction

As computer technology has advanced, information processing in command, control, communications, computers, intelligence, surveillance and reconnaissance (C4ISR) systems has become highly complex. The information processed by these systems is usually of a very highly sensitive nature and is often sensitive in both hierarchical (top secret, secret, etc.) levels and non-hierarchical (CANUS, CANUK, AUSCANUK, etc.) levels. Often, this information is entered into specific systems that are physically isolated from each other so that mandatory access controls can be maintained. The physical isolation of these systems makes it cumbersome to exchange information between systems of overlapping security policy. The result is inefficient sharing of sensitive information in situations where timeliness of exchange could be a life or death reality.

Since the mid 1990's, increasing efforts have been placed on improving coalition operations. Many systems have been created with the goal to improve the sharing of information and collaborative planning across coalition boundaries. The usability of these systems have had mixed levels of success.

In general, information domains are separated based on sensitivity of the information that is processed within the domain. In some cases, such as in the case of multi-level secure (MLS) systems, information of differing sensitivity may be processed within a single system. However, in the case of MLS systems, the information is still contained by the technology to prevent the inadvertent release of information from a higher security domain to a lower security domain. In each case, most systems still follow the Bell-LaPadula security policy model [1] for control of access to information of particular sensitivity levels. Basically, the Bell-LaPadula model allows access to information objects based on a "write up" and "read-down" policy. This means that a subject at a lower sensitivity level can write into an equivalent or higher sensitivity level, while a subject of a higher sensitivity level can have read access to information of an equivalent or lower level sensitivity level.

While the Bell-LaPadula model is a key policy model for information security, considerations beyond strict hierarchy of information causes complications. For example, many national military systems maintain additional caveats on information that are not strictly hierarchical. These caveats may be "eyes only" caveats like CANUS, CANUK, etc. – or particular operational codewords. Another good example is the NATO caveat that is placed on information generated within that information domain. These caveats often create subsets of information sensitivity (sometimes termed "non-hierarchical") equivalence that become complicated to control in a coalition environment.

Furthermore, regardless of the presence or absence of particular sensitivity levels, each nation within a coalition has national sovereignty considerations that must be handled in a coalition environment. In most cases, each nation connecting to a coalition information domain will place a firewall between their national systems and the coalition domain. The role of the firewall is to establish access and information flow control between information domains. In addition to firewalls, depending on the nature of the information, additional cryptographic mechanisms will be used to ensure the confidentiality of information in transit.

In addition, cryptographic mechanisms provide the means to verify the integrity of information when received. While firewalls and cryptography systems provide significant measure of control over access to and exchange of information, they create a complex set of intermediary systems between two operational users. A basic coalition connectivity picture is provided in Fig. 1. This diagram represents the conceptual connectivity between any nationally

sovereign operations environment to a coalition operational environment.

In the diagram in Fig. 1 the national system operations represent systems that are under the sole control of a single participating nation. These would be systems that are within the sovereignty of a particular nation. In the coalition operation, there may be many such national systems connected to the coalition operations domain. The coalition operations domain is a shared domain of many participating nations and may often be created for specific operational purposes.

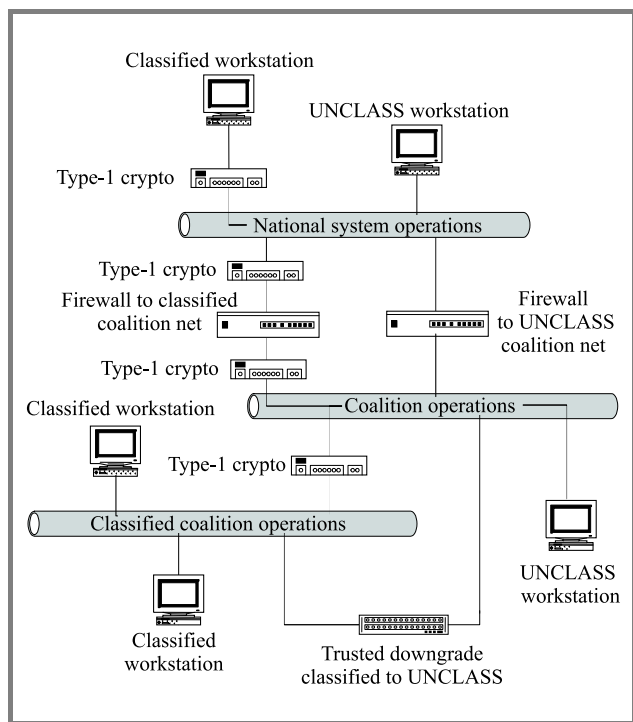


Fig. 1. Basic coalition connectivity.

Throughout the remainder of the paper, this basic coalition connectivity diagram will be used to highlight the issues related to operational needs, created as a result of the presence of intermediary systems. For the issues described, this paper will also describe three advances in telecommunications technology that could be leveraged to significantly improve current coalition operations. These technologies are:

- use of session border controller (SBC) technology to improve interworking of collaborative planning tools across coalition boundaries;
- use of advanced pattern matching engines for improved information sharing across national domain boundaries;
- use of multi-protocol label switching (MPLS) explicit routes (ER) to monitor/control packet flow of sensitive traffic.

2. Collaborative planning tools

2.1. Connectivity and protocol issues

Referring to Fig. 1, systems operating in the UNCLASS national operational domain that are connected to the coalition operational domain must pass traffic through a firewall. Note that these information domains may alternatively be at some other operationally equivalent sensitivity level (i.e., sensitive but unclass (SBU)). The firewall controls access and information flow between these information domains. In general, the firewall policy will be set, such that all traffic is blocked, except particular types of traffic between particular individuals or machines. In general, the policy will, as a minimum, place controls based on a 5-tuple (IP source, IP destination, TCP source, TCP destination, protocol type). Beyond the policies established by the firewall, network address translation (NAT) compounds the problem due to the need for mapping and manipulation of IP addresses at the information domain boundaries.

Classified workstations operating in the national system domain are connected in a similar manner via a firewall. However, note also that Type-1 cryptographic systems are used to protect the information as it passes through the UNCLASS (or other lower operational domains). As a result, the traffic must be decrypted prior to the firewall in order for the firewall to take appropriate policy decisions and then be re-encrypted in the coalition environment until it reaches the classified coalition domain.

In general, security policies are established such that connections between national systems and coalition systems must be initiated from within the national system and must use well known (pre-defined) ports. Any unused port is explicitly blocked. An example of such a policy has been used by the Canadian Forces in the Joint Warrior Interoperability Demonstrations (JWID) and has been described in [2].

The strict nature of such a policy ensures a strong measure of control over the flow of information; however implementations limit capabilities for true collaborative planning. Many collaborative planning tools use a signaling channel, on a known port, to negotiate a data or session channel port. Since the data or session port is not known a priori, the firewall policy is usually configured to block such traffic. The alternative is to leave a block of ports open such that the negotiated data port is allowed. While leaving these ports open solves an operational issue, it also creates additional security risk. This additional risk is usually deemed unacceptable and therefore the ports are closed and the application is denied.

Collaborative planning tools that fall into this category include voice over IP (VoIP), whiteboarding, chat, video teleconferencing (VTC), instant messaging (IM), etc. All of these applications serve a significant role in coalition collaborative planning and most remain a technology challenge to allow the connection of these tools to national systems. Thus, the utility of the collaborative planning tools is limited when national systems specifically deny connection.

While it may be expected that these limitation have been solved in recent years since publication of [2], this is not the case. Since the beginning of 2004, the US government has released at least two separate solicitations regarding the issues surrounding collaborative planning tools [3, 4]. It is recognized that “efficient, seamless ways to share information of varying classification levels and political sensitivities over a single network do not currently exist” [4].

2.2. Session border controller

In 2003, SBCs were voted the number 1 hottest new technology by Telecom Magazine [5]. Used in the telecommunications industry, SBCs exist “to provide a demarcation point between two service providers’ VoIP networks, allowing them to manage signalling and control routing for VoIP traffic” [5]. The key ideas behind the SBC are the management of signalling traffic and the routing of the data traffic. Originally designed to facilitate call setup for VoIP traffic, the concepts behind the SBC create a controlled interface between two network domains that are ideally suited for multi-media applications and protocols – the types of protocols that support coalition collaborative planning tools.

Sitting at the interface between two information domains, the SBC intercepts the signalling protocol between two systems. Taking the example of VoIP, the SBC will monitor for either of two dominant standards, the session initiation protocol (SIP) or H.323. When intercepted, the SBC either directly manages and controls the connection using an internal application level gateway (ALG), or it uses a separate control interface protocol to communicate with an external ALG system. The mechanisms required to support the external ALG system are being defined by the Middlebox Communications Working Group (MIDCOM WG) of the Internet Engineering Task Force (IETF).

While the details of SIP and H.323 differ significantly, the goal is the same, to establish a voice connection between two VoIP end points (phones). This paper will focus on SIP as “some observers believe that SIP will become dominant” [6]. SIP is “an application-layer control (signalling) protocol for creating, modifying, and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution, and multimedia conferences” [7]. One of the key benefits of SIP in the context of coalition interoperability is the fact that it is media independent. This means that the SIP itself is not tied to a particular media type (i.e., voice), but can be used for virtually any type of media traffic (i.e., video, instant messaging, whiteboarding). This is because negotiation of the media type and the parameters of the session are negotiated during the call setup process.

SIP supports five facets of establishing and terminating multimedia communications [7]:

- user location: determination of the end system to be used for communication;
- user availability: determination of the willingness of the called party to engage in communications;

- user capabilities: determination of the media and media parameters to be used;
- session setup: “ringing”, establishment of session parameters at both called and calling party;
- session management: including transfer and termination of sessions, modifying session parameters, and invoking services.

A simplified illustration of SIP being used to initiate and control a session between two end points is provided in Fig. 2. The initiating end point sends an “Invite” request to the recipient. Among other fields, the “Invite” will contain several fields that relate to the routing of the call to the recipient. These fields are the “Via” containing the address expected for response; “To” containing the destination universal resource identifier (URI) and “From” field containing the source URI.

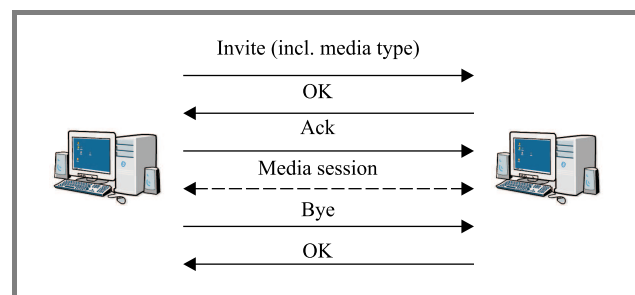


Fig. 2. Session initiation protocol.

In addition, the SIP invite will contain a “content-type” field that is used to identify the media application type that will be described in the body of the SIP invite message. For example, the content-type field may identify “application/SDP” to identify the session description protocol (SDP), used for voice. The body of the SIP invite would contain parameters associated with the SDP that could be processed by the recipient to enable the media type. The “content-type” and the body of the SIP payload are critical to the SIP message as it will contain the IP address and dynamic port assignments for any collaborative planning applications. This information is essential for the operation of the session border controller, as will be described below. Assuming that the parameters of the invite are accepted, and then the recipient of the invite responds with an “OK” message. This message is acknowledged by the originator using an “Ack” and the media session is established. During the course of a session, either end may add or remove other media sessions or types as required, by negotiating these sessions through the existing SIP session. When either end point terminates the session, a “Bye” message is sent to the other party, which is acknowledged by an “OK” message.

Not included in the illustration of SIP in Fig. 2 is the SIP use of proxies. Within the definition of SIP, the protocol allows for the use of intermediate proxies that are used to

relay the messages from the initiator to the recipient. Typically a SIP call will not go directly between two end points but will instead pass through proxies at the boundary between different network domains. In the case of a connection to a coalition network, a proxy would sit at the boundary between the national domain and the coalition domain. It is possible; however, that additional proxies will be required as the call request is passed through various points in the network. This would depend on the overall network connectivity.

A session border controller acts as a firewall that uses an application level gateway programmed to understand SIP. ALGs go deep into the data in the SIP packet and parse the “content-type” and payload. This allows the ALG to determine the IP addresses and dynamic ports that are required to enable the data ports of the collaborative planning applications. By understanding which ports need opening, the SBC dynamically opens only those ports needed by the application, leaving all others securely closed. This technique of opening small numbers of ports in the firewall dynamically is called “pinholing”. One of the key advantages of the ALG is that the constant monitoring of the session ensures immediate knowledge of call termination, allowing the “pinhole” to be closed immediately as well.

As described earlier, NAT causes difficulties in the use of collaborative planning tools. A SIP proxy is used to provide NAT traversal. The proxy has knowledge of the IP domain on both sides of the proxy, and separates the SIP call into two separate calls: one from the end point in the national domain to the proxy and one from the proxy to the coalition domain. The proxy is an intermediary control point and resolves the NAT issue. In most cases, the ALG will also incorporate a proxy and therefore is able to handle both NAT issues and the dynamic assignment of ports.

In some instances, the ALG of an SBC will be implemented in a separate device from the firewall. In this case, the SIP messages will be routed to the separate ALG for processing. The separate device will then dynamically control the firewall by telling it the IP address and UDP (or TCP) port information determined from the SIP payload. This approach using a separate device is being promoted by the MIDCOM WG in the IETF and is illustrated in Fig. 3.

The advantage of the MIDCOM scheme is that the firewall is not burdened by the impact of processing the SIP messages. Once the session is established, and pinholes are created, only the media streams are processed by the firewall and the impact of media dependent characteristics for latency, jitter and quality of service are minimized. In addition, once implemented with a MIDCOM interface, the firewall no longer needs to be upgraded for each new application service. Instead, the ALG can be upgraded separately, thereby minimizing the operational impact on the firewall.

In summary, the SBC provides a dynamic firewall solution that can be used to improve coalition operations. The na-




Fig. 3. ALG as separate signaling proxy.

ture of the connectivity to a coalition network creates difficulties due to NAT and the firewall policies that block applications using dynamic port assignments. The SBC uses proxy enabled ALGs programmed to understand SIP that dynamically open ports in the firewall. This dynamic control of ports is termed “pinholing”. The ALG may be implemented directly on the firewall, or it may be a separate device that communicates with the firewall using a MIDCOM scheme. The use of the ALG could allow collaborative planning tools to span the national and coalition information domain while minimizing the risks previously associated with the dynamic port assignments.

3. Trusted downgrade systems

Another area of connectivity that causes difficulties to coalition operations is the interface between higher and lower level sensitivity domains. As shown in Fig. 1, a trusted downgrade system would sit at the interface between the classified operational domain and the UNCLASS or SBU domain. The role of a trusted downgrade system is to allow the approved release of information from the higher level domain to the lower level domain. Recalling the Bell-LaPadula model, it is specifically denied by policy to “write-down”, that is, to pass information from the higher domain to the lower domain is forbidden.

However, this security model does not account for the “system high” nature of the classified operational domain. “System high” refers to the concept that all information in an information domain is treated as of the highest classification of information processed on the system. This means that despite the fact that some information may not be SECRET, if it is contained in a system that operates at the SECRET level then it is treated as if it is SECRET. Consequently, strict enforcement of the Bell-LaPadula model prevents the valid transfer of information from a higher domain to a lower information domain.

Recognizing this limitation, there have been several products developed and approved for operation that perform a trusted downgrade operation. The Radiant Mercury [8] is one such system developed by Lockheed Martin that has been in service at least since the mid 1990's. There are other similar systems in operation such as the ISSE Guard [9]. In general, these systems operate on a trusted computing base (TCB) that has been evaluated and approved in accordance with one of the trusted computing evaluation programs, such as the common criteria for information technology security evaluation (published as ISO standard 15405). As well, these systems enforce the control of information release from high to low in one of two manners. Either they use automated methods to approve release based on a review of highly formatted messages such as USMTF messages, or they rely on the approval of a release authority (i.e., approved email from a valid user with appropriate rights for release). A concise list of current MLS systems in use, including systems used for trusted downgrade is available at [10].

Taking the RM as an example, the RM release 3.0 serves two main roles and operates on a Sun platform. The RM has the capability to automatically review and approve the release of highly formatted message according to pre-defined rules. Due to the highly formatted nature of the message and the extensive rule base, it is possible for effective controls to be established such that only information appropriate to the lower sensitivity domain is released to that lower domain. The second feature of the RM is to approve the release of imagery files. In this case, the Radiant Mercury operates in a manner similar to most trusted downgrade systems; an authorized release authority must identify the file as approved for release and pass the file to the RM. The RM examines a header that has been applied to the file by the release authority and passes the approved image into the lower security domain.

The RM and similar products operate on a TCB that is based on a trusted operating system running on a general purpose processor. The operating system is responsible for the secure containment of information on both the high and low side of the system and is also responsible for the trusted transfer of information from the high side to the low side. This architecture has proved very useful, secure and is able to meet the needs of some operational requirements. However, with the growing scope of coalition operations, these systems may not be able to handle the increased demands placed on them. The architecture used forces a complete reassembly of the information content in order for the application to scan, parse, review, modify, approve and release all messages. This is a highly processor and memory intensive process that can be impacted by increased demands on the system.

On the other hand, modern networking equipment uses datapath technology that is designed to scan, parse, review and modify information at the packet level. For example, as a packet is received by a router, the packet header information must be scanned and parsed to extract the common

5-tuple information (IP source, IP destination, TCP source, TCP destination, protocol port). This information is used as a lookup key for access into a forwarding information base (FIB). Effectively, the FIB provides instructions to the datapath to inform it of the actions that should be taken on the packet. Often, these actions may include the modification of the packet (i.e., in the case of NAT) prior to forwarding to the release interface.

At first glance, the technology used in the switch and router datapath may seem ideally suited to the problem of trusted downgrade operations at very high rates. However, the technology widely used in switches and routers has been engineered for the specific problem of fast header inspection and forwarding. On the other hand, they are not well suited to scanning and parsing of information deep within the payload of a packet or where information spans across multiple packets. However, the basic technology used in switches and routers has formed the basis for advanced pattern matching engines (PME) that support inspection at higher layers.

3.1. Pattern matching engines

Pattern matching engines have been developed to support a variety of applications. These applications range from content switching to intrusion detection and prevention systems (IDS/IPS) to automated anti-virus platforms. PMEs have also been called content inspection engines (CIE) or simply search engines (SE).

There are two main types of PMEs, those that are based on a Ternary Content Addressable Memory (TCAM) technology and those that are algorithmic based. PMEs may optionally support both exact match and regular expression matching criteria. The TCAM type of PME has a direct relationship to the technology in use in the switch and router datapath for packet forwarding.

PMEs provide considerable potential for the improvement of the capability of trusted downgrade systems. The matching engine provides the baseline technology needed to perform high speed parsing and review of received information and the technology is intended to provide pattern matching capability deep into packet contents and across packet boundaries. An anti-virus scanning feature is a typical example where the entire payload of a data stream may need to be scanned. These applications exist in current products available on the market.

In some cases, PMEs are available directly from component suppliers. The PAX.portTM line of devices from IDT [11] is an example of PMEs available from a component supplier. According to the Linley Group, IDT, Inc. is currently the market leader in SE components, with Cypress ranked second [12]. PME components hold promise for the future of trusted downgrade platforms since they offer a flexible program matching language that can support a wide range of applications. Unfortunately, this also remains their biggest challenge. Following development of a system, considerable effort may be required to program the device for particular application needs of trusted down-

grade applications. Also, the technology is generally used for matching against relatively short “keys” or “strings”. Additional research is required to determine the suitability of such devices for the multiple match criteria that may be required, for example, in applications supporting military message formats (i.e., USMTF).

In addition to the general device supplier category, there are system vendors that have performed extensive research into PME technology and use this technology in their platforms. In this case, the technology is often tuned to the particular application space. For example, the FortiGate™ line of products from FortiNet™ makes use of the FortiASIC™ to perform fast pattern matching for anti-virus applications [13].

At this stage, additional research into the use of PME technology to support trusted downgrade applications is required. However, it appears that this technology could provide better baseline platform capability than the general purpose processors architectures in use today.

4. Controlled flow of sensitive traffic

4.1. Nature of provider networks

In a multinational coalition operation, connectivity for both tactical and strategic networks is established through network paths that are likely not a part of the normal grid used by these nations. As new network paths are created, it is often a requirement to lease the network from providers. This leaves the network connectivity paths outside the control of the national military force requesting the service. This creates a situation where the provider may route the traffic through other nations where the owner of the data may not want traffic to pass.

In fact, even using the standard strategic networks that provide for normal national operations, the provider may route traffic through areas where the national military may not want the traffic to go. Fortunately, in the case of normal day-to-day operations, the provider provisioned network is often controlled by strict contractual agreements that preclude routing of traffic through particular areas of the world.

Before considering the mechanisms available to control the flow of sensitive information, it is important to highlight why the traffic routes are a concern. On the one hand, all sensitive operational traffic will be protected from disclosure by some cryptographic means. In the case of classified operational traffic, Type-I cryptography is used, thereby providing the assurance that even if intercepted, the traffic is unreadable and is therefore protected. This being the case, there it can be argued that the route taken by the traffic is not a concern.

On the other hand, information system security needs to account for the integrity, availability and accountability of the information, not just the confidentiality. Integrity protection ensures that any modification of the traffic is detected. In addition, it is desirable that opportunities to modify the traffic be minimized. It may be beneficial to know that

the traffic has been tampered with, but this doesn't help with the fact that the correct data has not been received. Also, availability concerns highlight the importance of ensuring that there are no interruptions to service guarantees from the provider.

Given these concerns, an argument can be made that national bodies may still desire to have added control of the path that traffic takes within a provider network. The nature of routed data networks does not really support this type of control. Routing protocols are designed to negotiate best path options for traffic. While the network provider can establish controls of the paths, the path options are based on the provider considerations for best path, not the considerations of the data owner. For example, a provider will establish paths to maximize bandwidth utilization and meet quality of service (QoS) guarantees. In some cases, this may result in the passing of traffic over links that reside in hostile locations. It would be beneficial for the coalition operations partners to have some measure of control over the approval of traffic.

4.2. Multiprotocol label switching and explicit routes

In traditional routed networks, the routing of traffic is based on the address of the destination of the packets. In the case of most networks today, this address is an Internet Protocol (IP) address. By contrast, in label switching, “instead of a destination address being used to make the routing decision, a number (a label) is associated with the packet... a label is placed in a packet header and is used in place of an address (an IP address usually), and the label is used to direct traffic to its destination” [14].

Label switching provides several advantages to the network provider [14]; speed and delay, scalability, simplicity, resource consumption, and route control.

Route control is the key consideration in the context of control over coalition traffic paths. Route control allows the system to designate a specific route path from among many that may lead to the same destination. This is sort of like placing an “Air Mail” label on a letter. With the “Air Mail” label, the letter will take a non-standard path that, one would hope, has an improved delivery time. In the same way, a label can be used in a system to control the route taken. The provider can engineer the network to route high priority traffic to one set of resources, while lower priority traffic takes a different path. The removal of lower priority traffic from the high priority resources reduces congestion and ensures guaranteed service levels can be met.

Multiprotocol label switching is published under RFC3031 [15]. MPLS combines label swapping and forwarding with network layer routing. “The idea of MPLS is to improve the performance of network layer routing and the scalability of the network layer” [14]. Within MPLS, a label switch path (LSP) is established either through a route negotiation protocol (i.e., link determination protocol), or through constraint-based routing. In constraint-based routing, the LSP is established manually.

It is possible to combine both automated route protocol establishment and constraint-based LSP configuration. In this case, the automated routes would be restricted by the configured constraints. These constraints are often associated with quality of service. A router that is aware of MPLS is termed a label switch router (LSR).

There are two main methods that MPLS uses to choose an LSP between nodes. In the first method, the LSR is free to independently select the next hop LSR based on knowledge it has in its routing table. The second method is called explicit routing. ER is used to define constraints on the LSP by identifying specific LSRs that must be used in the LSP. Assuming that provider networks are MPLS capable, the concept of ERs could be extended to define constraints on the LSP that prevent coalition traffic from passing through routers that reside on undesired traffic paths.

Note that to control traffic based on sensitivity, MPLS would need to be extended. As described earlier, ER is generally used to provide QoS guarantees. Within the various methods to negotiate an LSP, there is no real concept of LSR location information, nor is there any notion of an "approval to process" identifier. This information would be critical to the extended use required to control the flow of sensitive information. Furthermore, the constraint-based link determination protocol (CR-LDP) used to establish LSPs would need to be extended to include an authentication mechanism that includes both location and approval to process criteria.

5. Evaluation considerations

The methods described in this paper are not specifically related to security products. However, the use of these methods to support coalition interoperability is identified to ease the burden of constraints placed on coalition operations due to security concerns. Therefore, it is important to note that incorporating these methods into systems supporting coalition interoperability will require trusted product evaluations and certification and accreditation for operation. Current research into the use of these technologies has not included any consideration for product security evaluation, though there is no known reason to believe that these techniques could not be included in a secure system design.

6. Conclusions

This paper has examined several issues related to coalition interoperability. These issues related to:

- the denial of collaborative planning tools across the national to coalition boundary;
- processing requirements for trusted downgrade platforms;
- the controlled flow of sensitive traffic.

Despite development and operational deployment of many systems, each of these topics remains a challenge to coalition interoperability.

This paper has identified three technology advances that could be used to improve coalition interoperability. These technology advances are:

- the session border controller;
- advances in pattern matching technology;
- use of multi-protocol label switching explicit routes.

Integration of any of these technologies will require trusted product evaluation and certification and accreditation for operational approval.

References

- [1] D. E. Bell and L. J. Lapadula, "Secure Computer Systems: Unified Exposition and Multics Interpretation", MTR-2997, Rev. 1, MITRE Corp., Bedford Mass., March 1976.
- [2] D. Wiemer, "Wiemer-Murray domain security policy model for international interoperability", in *Proc. NISSC*, Arlington, USA, 1998, <http://csrc.nist.gov/nissc/1998/proceedings/paperF20.pdf>
- [3] Defense Information Technology Contracting Office (DITCO), "DRAFT Next Generation Collaboration Service, Statement of Objectives", published as pre-solicitation, 27 Jan. 2004.
- [4] Department of the Navy, Naval Supply Systems Command, "Multi-national Information Sharing Environment", published as pre-solicitation notice, no. H1967, 23 Jun. 2004, <http://www1.eps.gov/spg>.
- [5] Telecommunications Magazine, "10 hottest technologies", Apr. 2003, <http://www.telecommagazine.com/default>
- [6] National Institute of Standards and Technology, Special Publ. 800-58, "Security Considerations for Voice Over IP Systems", Apr. 2004, http://csrc.nist.gov/publications/drafts/NIST_SP
- [7] J. Rosenberg *et al.*, "SIP: Session Initiation Protocol", RFC3261, Internet Engineering Task Force (IETF), June 2002, <http://www.ietf.org/rfc/>
- [8] J. Pike, "FAS intelligence resource program: Radiant Mercury (RM)", Jan. 2000, http://www.fas.org/irp/program/disseminate/radiant_mercury.htm
- [9] ISSE Guard Program Office, "ISSE Guard, suite of products for secure multi-domain information exchange", Nov. 2003, <http://www.rl.af.mil/tech/programs/isse/>
- [10] D. Zellmer, "Multi-level security: reality or myth", As part of GIAC practical repository, SANS Institute, March 2003, http://www.giac.org/practical/GSEC/Douglas_Zellmer.pdf
- [11] Integrated Device Technology, Inc., "IDT™ PAX.port™ 1200 Content Inspection Engine", 2003, http://www.idt.com/docs/76T1200BH_BR_88384.pdf
- [12] The Linley Group, "Search engine market maturing", *The Linley Wire*, vol. 4, issue 11, June 2004, <http://www.linleygroup.com/npu/newsletter/wire061004.html#2>
- [13] FortiNet™, "Comprehensive solutions for real time network protection", 2004, <http://www.fortinet.com/doc/FortinetBroch.pdf>
- [14] U. Black, "MPLS and Label Switching Networks". Prentice Hall, 2002.
- [15] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture", RFC3031, IETF, Jan. 2001, <http://www.ietf.org/rfc/>



Douglas Wiemer is a retired Canadian Armed Forces Captain. While serving, he was employed in the Air Force as a Communications and Electronics Engineer and specialized in information security on data networks for Command, Control and Intelligence Systems (CCIS). Among other roles, he served as the System Security Engineer for the planning of the Canadian participation in the Joint Warrior Interoperability Demonstra-

tion 1997 (JWID'97). Douglas Wiemer is currently employed in the Alcatel Networks, Research and Innovation (R&I) group on the Intelligent Switch Routers (ISR) project. His research is focused on architecture studies of the switch and router datapath. These studies are used to assess the capabilities and technology needed in the datapath to enable advanced services like the Session Initiation Protocol (SIP).

e-mail: douglas.wiemer@alcatel.com

Alcatel Networks

Research and Innovation Group (R&I)

Intelligent Switch Routers Project (ISR)

600 March Rd, Ottawa, Ontario, Canada