Improving Security in Coalition Tactical Environments Using an SDN Approach

11

Vinod K. Mishra, Dinesh C. Verma, and Christopher Williams

11.1 Introduction

The principles of software-defined networking (SDN) can be used in many types of networking environments and in particular for addressing challenges related to security and operations of these networks. The design philosophy embedded within SDN has applicability far beyond the control of data center networks where it first originated. However, in order to address the unique requirements of new environments, the implementation and architecture of SDN need to be modified suitably so that it can be applied in an efficient manner to them.

In this chapter, we explore how SDN can be used to improve the security and situational awareness in tactical networks in general and in coalition tactical networks (CTN) in particular. We assume that the reader is already familiar with the principles of SDN, which have been elaborated upon in other chapters of this book. However, tactical networks and coalition tactical networks are special type of networks which the reader may not be familiar with. Therefore, we first introduce these two types of networks with emphasis on their special security and networking requirements. After an introduction to the networks, we will look at the ways in which the SDN concepts can be applied to them to address their networking and security requirement.

V.K. Mishra
U.S. Army Research Labs, APG, Aberdeen, MD, 21005, USA

D.C. Verma (✉)
IBM T J Watson Research Center, Yorktown Heights, NY, 10598, USA
e-mail: dverma@us.ibm.com

C. Williams
Defence Science and Technology Laboratories, Porton Down, Salisbury, Wiltshire, SP4 0JQ, UK

S.Y. Zhu et al. (eds.), *Guide to Security in SDN and NFV*, Computer Communications and Networks, DOI 10.1007/978-3-319-64653-4_11

As we apply the concepts of SDN to CTN, we adopt the approach that is the defining feature of the SDN architecture, namely, the separation of data plane (DP) from the control plane (CP), and consolidate the CP functionality at a central location in the network. The function of any computer communication network is to accept data packets (also known as protocol data units) from one computer and deliver it to another computer. The network consists of several elements, each of which performs some operations on receiving the packet, e.g., (1) deciding which of several possible outbound interfaces to choose for forwarding the packet or (2) whether to drop the packet due to a security reason. These per-packet operations are the DP functions. In order to perform them, it is necessary to complete some operations earlier, e.g., (1) populating the entries within data forwarding tables, (2) setting up virtual connections, or (3) defining any packet filtering rules. These types of operations are the CP operations. In a traditional network, both DP and CP functions are carried out using a distributed algorithm; e.g., in an Ethernet, a forwarding table which follows the links of a spanning tree among all participating switches is established as part of the CP using a distributed protocol implemented within each switch. This results in both CP and DP functions residing on the same network device and is the ultimate reason for the inflexibility of non-SDN networks.

In SDN, the CP operations of individual network devices are replaced with CP operations run from a centralized SDN controller (SDNC). The SDNC implements the control plane operations as software running on standard IT servers. This moves the CP functions from each device in the network to a logically centralized controller and enables more flexibility. The high-speed DP that is responsible for actually forwarding packets remains in the network devices. As an example, in an Ethernet, the logically centralized controller can be configured to implement algorithms that compute not just a spanning tree but a more complex graph for forwarding packets which use links not on the spanning tree. Thus, it is not necessary to implement a distributed protocol, which is more complex, and may require standardization among different devices manufacturers to work properly.

Another key component of SDN is a set of programming interfaces called northbound interface (NBI) and southbound interface (SBI). NBI allows network applications and policy commands to be communicated to the SDNC. Similarly SBI is used by SDNC to control the data plane in network devices like switches and routers. These interfaces allow applications and control programs to automate network operations through well-defined, open APIs enabling much more agile interaction with the network than traditional methods, such as scripted command line interfaces (CLIs) and proprietary interfaces. Currently OpenFlow [1] is the dominant open SBI interface. It has been standardized by the Open Networking Foundation (ONF). There is at present no universally accepted open NBI, but attempts to define one are continuing.

For reasons discussed later in the chapter, having an SDNC as a physically centralized entity is not a good solution for tactical networks. Nevertheless, many benefits can be obtained by separating CP and DP functions in a CTN. The logically centralized SDNC can be also implemented in a physically distributed manner to improve its security and resilience.

11.2 Tactical Environments

In both military and civilian contexts, there are several situations when a group of people need to perform a task in an area where there may not be adequate infrastructure for communications. As an example, a platoon of soldiers or policemen may be asked to surround and secure a building in which suspected insurgents may be hiding. Similarly, a group of firemen may be dispatched to handle a fire in a mine or a forest where they may be outside cellular communication coverage. These environments, where a group needs communication without reliance on an existing infrastructure for a limited period of intense activity, are referred to as tactical environments.

In the tactical environment, the people who are involved in any operation would have mobile devices with them. Depending on the technology available to the group, they may be using autonomous vehicles like mules or drones to perform their tactical mission, and they may also be carrying equipment with built-in smart communications capabilities. Also network nodes themselves may move erratically; e.g., firemen may have to beat a hasty retreat if a sudden conflagration occurs, or a soldier may have to make sudden movements to avoid enemy fire.

Within a tactical environment, the nodes carried on person may be supported by more powerful (off-body) support system nodes. Firemen may have a supporting fire truck, which can carry an access point connected to a satellite network. Similar type of supporting infrastructure may be available to soldiers through one or more vehicles being driven in the area of operations. Furthermore, if the operation is being conducted in an urban area, the personal devices can even have cellular connectivity, and thus nodes in the tactical environment may be connected to the infrastructure at least some of the time, if not always.

11.2.1 Segments of a Tactical Network

Although the details vary widely in different countries and military and civilian organizations, one can create a simplified and abstract model of the network of any organization with a tactical environment component as shown in Fig. 11.1.

The model shown presented above has four segments:

- The first segment consists of the devices and the network used by people at the very edge of the operation. This will contain many different handhelds; unmanned aerial vehicles (UAVs); intelligence, surveillance, and reconnaissance (ISR) devices; mobile networking; and computing environments to be carried onto various platforms such as tanks, ships, or vehicles. These devices may establish an ad hoc network among themselves or use satellite communications to interconnect themselves. Thus, the tactical environment would be a mobile ad hoc network (MANET), but instead of being completely ad hoc, the tactical

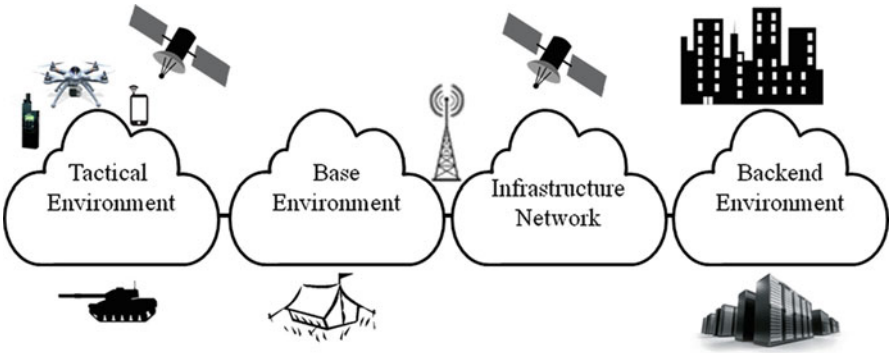

Fig. 11.1 Simplified model of a single organization network

environment can usually rely on a limited amount of infrastructure support from other components in the single organization network.

- The second segment connects tactical environment to the base environment found in bases and buildings used as support operations. In a military context, base environments may use portable laptops, desktop computers, storage devices, and networking equipment that usually create a temporary network infrastructure. It may also reflect the computing environment in a base that has been set up for a temporary period, ranging from a few days to a few months. The base is set up to provide logistics support to soldiers who may need to operate over a large area. In a civilian context, the base may be a temporary camp setup for operations in a disaster recovery area; e.g., it may be the place where people affected by a flood can find shelter, or it may be the command center of teams of firefighters trying to put down a forest blaze.
- The third segment consists of the infrastructure network which connects to the second segment of the base networks. It may contain satellite communications or may leverage installed infrastructure such as cellular communications networks. Depending on base's lifetime, it may also leverage fixed network infrastructure, such as a wired cable of fiber network to connect to the fourth segment of the backend environment.
- The fourth segment consists of the computing infrastructure found in buildings and military headquarters. In general, computing and communication resources in infrastructure and backend environments are plentiful and tend to be static in nature. It should be remembered that this bandwidth abundance becomes quite scarce by the time it reaches the tactical edge.

11.2.2 Security Considerations in Tactical Networks

One of the key attributes of a tactical network is that it is formed quickly and is decommissioned once the mission is completed. As a result, the security considerations in tactical networks are very different from those in the fixed infrastructure networks.

Many tactical networks need to operate in hostile environments. In the context of a fire, there may be a sudden conflagration, and due to the need to move suddenly, the firefighter tactical network may lose a node unexpectedly. In a military tactical network, enemy action may bring down a node at any time. A tactical network needs to be able to provide a high degree of resilience in the face of these sudden node failures. Adversarial action may also include less dramatic activities, such as an arsonist putting in a malicious node in the firefighter scenario or an enemy putting in malicious nodes to join the tactical network in an attempt to intercept and manipulate communications. Every tactical network needs an approach by which only authenticated and authorized nodes are able to join in the tactical environment. In addition, the communication needs to be encrypted so that they are not intercepted easily off the air.

The need for authentication needs to be balanced with the need to be able to form the network rapidly. This means that the process for authenticating nodes that can join the tactical network needs to be very agile and rapid. When a platoon of firefighters (or that of soldiers) is called upon to take on a mission, they cannot go through a complex manual process for establishing keys and certificates for authentication and encryption of communications. The process for establishing the required keys/certificates needs to be extremely agile and not add to the preparation time for the platoon members.

Another consideration in preparing the nodes that will constitute the tactical network is the fact that in a tactical network, all nodes may not be initially present at the formation time. New members in the tactical network may need to be added later. In the case of a platoon of firefighters, some firefighters may arrive late to the site of fire due to traffic delays. In the case of a military tactical network, additional troops may be added to a mission, and their nodes need to be able to join the existing nodes in the tactical network. Thus, the authentication mechanism needs to be able to support new authorized nodes while preventing access by unauthorized or malicious nodes at the same time.

The communication among nodes in a tactical environment would need to be encrypted, but the bandwidth capabilities and processing capabilities of the nodes are usually limited. As a result, the encryption mechanism needs to be a lightweight one. Either a shared key mechanism or popular TLS/SSL protocols can be used depending on the processing capabilities of the mobile node devices. These protocols use certificates and public key cryptography approaches to set up the shared keys that can be used for a brief period of time. In either case, the right shared keys or the right certificates including the public keys need to be set up in the nodes.

In addition to having a lightweight configuration, the nodes in the tactical environment also need to react to impending threats that they observe in the environment during the operations. These include having nodes react to sudden movements that are caused due to mission requirements, dealing with sudden loss of neighboring nodes, reacting to any intrusion attempts that an adversary may be making, and dealing with attempts of an adversary to jam communications. Each node needs to have the appropriate intelligence and insight to deal with these situations as they arise.

11.3 SDN Architecture for Tactical Environments

The SDN approach to networking was applied initially to the regular data center or fixed infrastructure IP and wireless networks. There are three key differences between such networks and tactical environments which need to be accounted for when using the SDN approach in tactical environments.

- *Bandwidth Constraints*: The bandwidth on wireless links is significantly lower than that on wired networks. Due to mobility, electromagnetic interference, and spectrum limitations, tactical network link's effective bandwidth is in the order of a few kilobits per second for platoon level links, and trunked wireless links can go up to a few megabits per second. This is significantly smaller than the available bandwidth in data center high-speed networks, where optical technology can easily offer hundreds of gigabits per second or higher bandwidth. Loss in wireless networks also tends to be high, from 1 to 10% [2], and this causes additional challenges for communication.
- *Disruptions and Failures*: Nodes in the tactical environment can fail suddenly, and due to the mobile nature of the network, a network node may not always have connectivity to the SDNC.
- *Short lifetime*: The lifetime of a traditional network is a few orders of magnitude larger than the time it takes to set up them or tear them down. In contrast, the lifetime of a tactical network is comparable to the time it takes to set up or tear down the network. As a result, the SDNC for a tactical network needs to take on additional functions, which may not be considered by the SDNC of a backend or data center network. Specifically, an SDNC for a tactical network needs to also handle the situation in which a node needs to join or leave the network.

Despite these challenges, SDNC architecture confers significant advantages in managing the security and operations of tactical environment. We propose such an architecture, in which the SDNC controller is part of the support infrastructure for the tactical environment. In the case of firefighters addressing a building fire, the SDNC may be a computer on a fire truck near the scene of the fire. In the case of a military operation, the SDNC could be a computer on a support vehicle. Nodes that join the tactical environment communicate for a brief period with the SDNC to get the appropriate security credentials for them to join the network. The SDNC can also provide policies to drive the operation of the nodes in the tactical network for routing and traffic control. Once a node is part of the network, it may or may not be connected to the SDNC. The SDNC needs to make provisions for enabling the network node to function properly even when it is not connected to the node. When nodes leave the network, they can briefly communicate with the SDNC to make a graceful exit from the network.

11.3.1 Challenges of SDN Architecture in Tactical Environments

The traditional SDN architecture requires a node to contact the SDNC whenever it encounters a situation in which it needs to make a decision for a DP operation for which it does not have the required CP information already within its control. In an environment like a data center where there are few losses and the bandwidth not constrained, having the node contact the SDNC for every decision is a nonissue. However, in wireless environments, where connectivity to SDNC cannot be assured, bandwidth is limited, and communication is lossy; contacting the SDNC for every CP decision is not a viable approach.

In order to deal with this situation, we propose an architecture in which the SDNC provides each node with the appropriate configuration and policies when the node first contacts it for access to the tactical environment. Configuration refers to any information that is needed by the software on the node to perform its operations and consists of files which include various parameters such as the type of security protocol to use and the size of encryption keys to negotiate for, along with any required security certificates or security keys. The policies provide a set of rules which tell the node how to react in different situations that may be encountered by the node and consist of actions to be taken under those conditions.

Using the model described above, the implicit assumption is that each node has an agent which invokes a standard interface to communicate with the SDNC. The agent considers its local configuration and policy data to be a cached copy of the information at the controller. The cached information can be maintained in sync with the information at the SDNC by using traditional cache coherency approaches, e.g., by having the agent check for any updates periodically with SDNC or by having the agent check that the configuration or policies have not changed when it needs to be invoked for a CP operation. Unlike a traditional node architecture, the agent can support not just policies that are common across nodes but also node-specific policies that are determined by the SDNC.

With this model, the traditional set of APIs defined in the SDN architecture get modified slightly as shown in Fig. 11.2. In the figure, the left-hand side shows the standard API definitions for the SDNC with an NBI exposed to the user/application to define the policies or configuration to the controller and a SBI that enables the communication of those policies and configuration to different network nodes or more specifically the DP component in each of the nodes. In the tactical environment, the device contains a DP as well as an agent that acts like a proxy for the SDNC. The SBI provides the interaction between the DP and the agent, the NBI remains unchanged in its functions, and an additional interface, the configuration coherency interface (CCI), is introduced between the nodes and the controller.

The CCI allows the DP elements to interact unmodified with the SDNC in coalition contexts, except that the interaction is now happening with the agent. In a tactical environment, the agent may frequently lose its network connectivity with the SDNC. The CCI deals with this loss of connectivity. In an environment when the device may only be connected with the SDNC occasionally, the agent

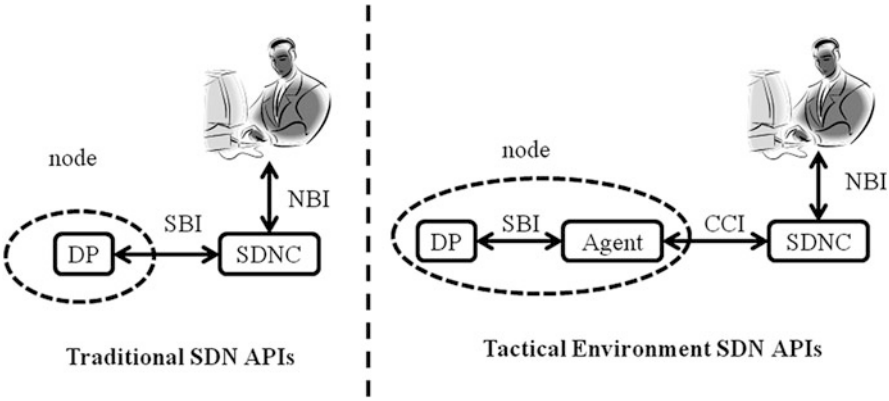

Fig. 11.2 APIs for SDN in tactical environments

can get the configuration and policies that are needed for the DP from the SDNC while connected. When it is disconnected, the agent is still capable of providing the required interaction between the DP, which is unaware that the real SDNC is disconnected. When the SDNC is reconnected, the agent can use the CCI to refresh its policies.

An obvious implication of the architecture for SDN in tactical environments is that the value of a standardized SBI in tactical environment is significantly less than that in traditional SDN environments. On the other hand, the need for defining a standard CCI for maintaining coherence between the agent and SDNC is important to enable devices and controller from different manufacturers to interoperate.

The other key difference in SDNC for tactical environment requires looking into the life cycle of a tactical network as well as that of nodes within the tactical network in more detail.

11.3.2 Life Cycle of a Tactical Network

The life cycle of a tactical environment can be described in a three-stage process as shown in Fig. 11.3.

These stages have specific functions:

1. In the first or the planning stage, the network functions like the type of authentication mechanism, encryption protocols, and policies for the operation of the network are defined.
2. In the second or the operation stage, the network is active, and it helps the performance of the mission.
3. In the third or the decommission phase, the mission of the network has been achieved, and the nodes are in the process of leaving the network.

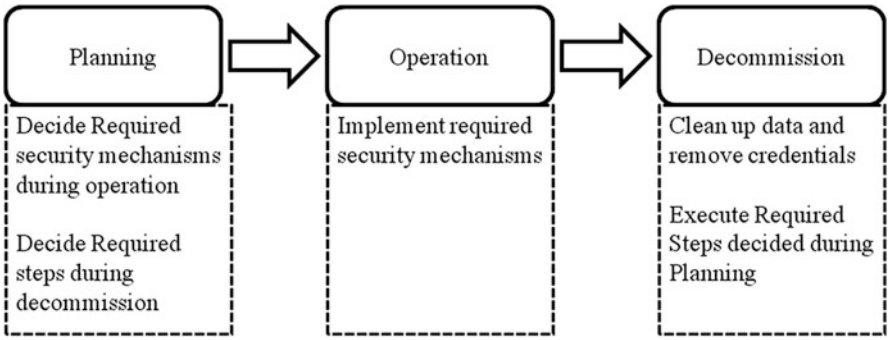

Fig. 11.3 Life cycle of a tactical environment

In each of these three stages, the SDNC needs to be able to provide the right control instruction for the different nodes in the network as the network progresses through each of these three stages. During the decommission stage, attention needs to be paid to the fact that some of the nodes may need to purge themselves of sensitive information if they are being taken out of service.

The policies and configurations needed for each of these stages is provided using the NBI to the controller. The NBI can either be invoked by a human administrator or by another software program. These policies and configuration will be provided to the agent in different network nodes as each of the network nodes goes through various states in its life cycle in the network.

11.3.3 The Tactical Network Node States

The tactical network nodes go through several states as they join the network during different life-cycle stages of the network. Some of the typical node states are shown in Fig. 11.4. In each of these states, the network node needs to determine what its configuration ought to be depending on the current network situation.

- Pre-authorized state: When a node has not been configured with the right credentials to join the tactical network, it is in the pre-authorized state.
- Authorized state: When that node talks to the SDN controller and gets the right keys and certificates to connect to the tactical network, it is in the authorized state. The node can then become operational when it joins the network and is connected to the other nodes in the network.
- Operational connected state: In this state, the node should typically have connectivity to the SDNC.
- Operational disconnected state: The node itself may be part of the tactical environment and connected to the network but still be unable to reach the SDNC. In that case, it is shown in an operationally disconnected state.

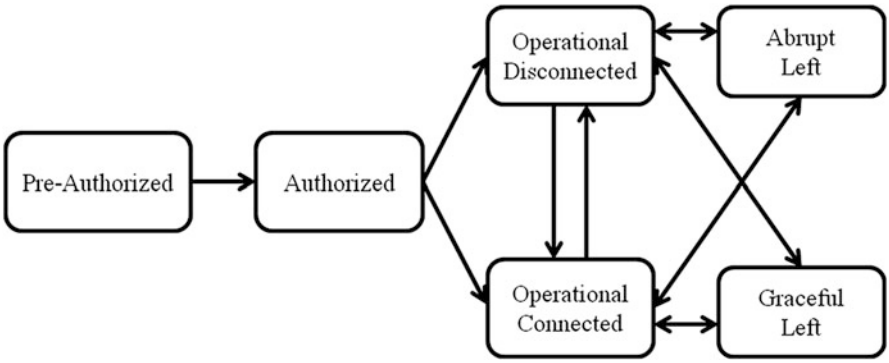

Fig. 11.4 States of a node in tactical environment

- Graceful left: After the mission in the tactical environment is over, the node may leave the network gracefully, i.e., after it communicates with the SDNC and performs the required cleanup procedures, if any.
- Abrupt left: On the other hand, some of the nodes may leave the network in an abrupt manner, e.g., leave the network without having this exchange with the SDNC.

When the SDNC is contacted by a node in the pre-authorized state, it needs to provide the node with all the configuration and policies needed for its operation in each of the stages of the life cycle, at least for those states where the node may not be connected to the SDNC. The SDNC may choose not to provide this configuration and policies for the node when it is able to connect to the SDNC, i.e., for the operational connected or graceful left states. However, it would be more bandwidth efficient for the SDNC to provide any configuration, policy, or other information needed for the control plane operations for those states as well.

When the node arrives in the pre-authorized state and contacts the SDNC, its credentials to access the network are validated, and on passing the validation checks, it is provided with the control plane information needed for operational disconnected state, as well as the abrupt left state. The operational connected state control plane information can also be provided. On receiving this information, the state changes to authorized. In this state, the node has the appropriate configuration allowing it to connect to the tactical network and communicate with other nodes in the network in a secure manner. Once it is operational, the node can check if it has a connection to the SDNC. It can then change its state to operational disconnected or operational connected state, appropriately.

In the disconnected or abrupt left state, the node uses the information provided by the SDNC during the authorized phase to perform its operations. In the connected state, the node uses the information provided in that phase as cached information. It can check if the set of configuration and policies have changed and use the locally cached information if it has not changed. It can download the changed control plane information if any update is available.

Table 11.1 Typical DP and CP functions at different life-cycle stages

Node state	DP function	CP information
Authorized	Joining the network	Security credentials
		Network configuration (e.g., SDNC address)
	Routing information	Routing configuration,
		Routing table
Operational	Forwarding	Forwarding table entries
	Filtering	Packet filtering rules/policies
	Encryption	Encryption keys/certificates
	Security monitoring	Intrusion detection/prevention policies
	Situational awareness	Situational awareness policies
Left	Reporting	Reporting policies
	Cleanup/retention	Data retention policies, data cleanup policies

The state of the network node defines the nature of the DP functions that need to be performed and the CP information to enable the DP functions. Table 11.1 shows some of the typical DP functions and associated CP information at different states of the network.

11.4 SDN-Based Operational Security and Situational Awareness

A tactical network node in the operational states (either disconnected or connected), as defined by the life cycle in Fig. 11.4, needs to obtain control plane information. The manner in which it is done is similar to that of SDN in backend networks, with the exception of caching that we described in the previous section. However, in order for the node to react to the security threats in the operational environment in an agile manner, it needs to be aware of its security situation and threats. In this section, we examine how situational awareness can be provided in a tactical environment.

11.4.1 The OODA Loop

In the human decision-making process, a common approach for situational awareness is the use of the OODA loop [3]. It explains that activity as consisting of the four stages of observe, orient, decide, and act as shown in Fig. 11.5.

The four phases are the following:

- Observe phase: All relevant information available from the environment is collected.
- Orient phase: The observation is analyzed further to get a deeper understanding of its implications; e.g., one may try to determine a root cause from the various observations.

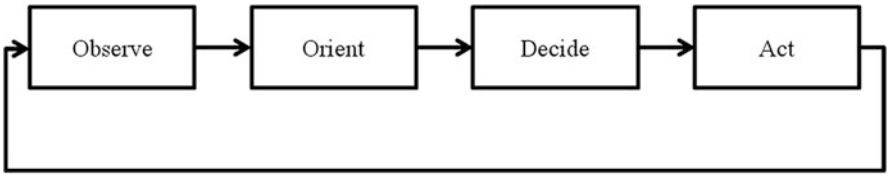

Fig. 11.5 The OODA loop for decision making

- Decide phase: The trade-offs involved in different courses of action are considered, and the right course of action is determined.
- Act phase: The action is actually undertaken, which results in a change to the environment, again leading to an observation of the environment. This completes the first loop. The next loop starts again with observing the changed environment.

As an example, suppose one hears the sound of a gunshot.

- During the observation phase, the sound of the gunshot is heard.
- In the orient phase, additional determination, e.g., the location of the gunshot, is determined, or other information sources, e.g., a camera video input, are used to get more information.
- In the decide phase, the possible options to deal with the gunshot is determined.
- In the act phase, the resulting action is then taken.

Although developed for the human behavior, the OODA loop can also be applied to tasks performed by a computer and in particular to the task of security in military networks. The application of SDN to cybersecurity situational awareness deals with using the OODA loop for cybersecurity to get humanlike situational awareness implemented within computer software.

11.4.2 Control and Data Plane Components of OODA Loop

In order to apply SDN principles, we need to differentiate between the control part and the data part of the cybersecurity situational awareness, as well as define what the implementation of the OODA loop means in this context. The architecture that we envision for cybersecurity situational awareness implements the OODA loop in software in the elements of environment which is described in Fig. 11.2.

It should be noted that there are actually two layers of OODA loops in operation: an outer or network-level OODA loop that directs the decisions and actions of the SDNC and an inner or device-level OODA loop that directs the operation of the DP with delegated authority (through policy) to act autonomously as directed by the SDN agent. The two-layered OODA loop architecture is shown in Fig. 11.6. The OODA loop software is responsible for performing the tasks required in the OODA loop as follows (for both device-level and network-level decision loops):

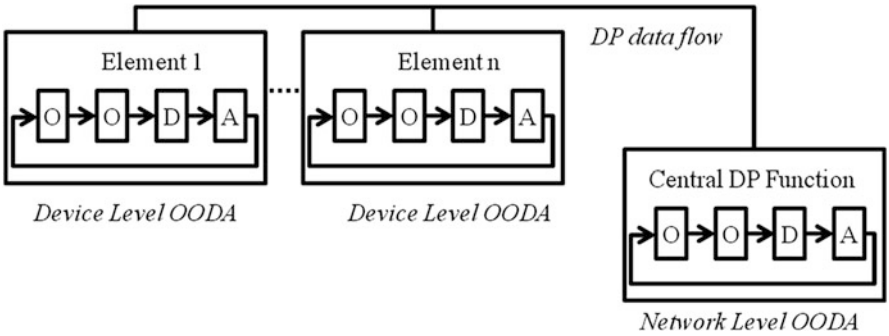

Fig. 11.6 Data plane with two layers of OODA loop

- The O (observe) part of the OODA loop consists of capturing portions of the network traffic that an element is seeing. Such data observations can be achieved by activating a variety of data collection elements. A data collection element may be collecting a subset of network packets, or looking at performance metrics within a computing system, or keeping track of the number of processes that are active within a computing device.
- The O (orient) part of the OODA loop needs to determine if anything abnormal is taking place in the network environment. The orient part of the system tries to map the observed data into higher-level phenomenon [4], implementing algorithms for root cause analysis to determine why specific observations might be happening. This may require the application of artificial intelligence (AI) and machine learning (ML) principles to discern patterns and classify the environmental cues.
- The D (decide) part of the OODA loop needs to process the collected information and assess the ongoing threat situation. On the basis of this assessment, a course of action (COA), which may comprise multiple parts, is then chosen as the response. Several approaches for making decisions, e.g., using policies or rules, utility maximization, and game theory, can be used at this stage. In the decide phase of the device-level OODA loop, one of the decisions may be to send the collected information or a subset of the information the network-level OODA loop for processing.
- The A (act) part of the OODA loop then implements the COA activities. For cybersecurity purposes, the action may consist of installing new network access control rules, information filtering policies, reconfiguration of security parameters, or switching to a different mode of encryption for secure communication, or specifically for the device-level OODA loop, the transmission of information to the network-level OODA loop. For the network-level OODA loop, one of the actions may be to update the policies and configuration of one or more device-level OODA loops

In each of the above implementations, the CP and the DP functions of the OODA loop can be defined, first for the device-level OODA loop:

- In the observe phase, the DP function is the actual collection of the data. Determining which type of information to collect and how to trade off the power and energy needs of an element against that of the normal computation would be the CP functions.
- In the orient phase, the DP function invokes the algorithms that map observations into phenomenon, while the CP function is the definition of the parameters in the algorithms that can enable such a mapping.
- In the decide phase, the policy rules (as defined by the CP from the network-level OODA loop) are implemented in the DP which leads to a decision. If a course of action cannot be identified to meet the required objectives within the policy constraints, then the DP can refer back to the SDNC for further guidance on the appropriate action to take.
- In the act phase, the DP functions then actuate the desired action of the data flow or sensors.

As an example, if a set of rules are being used to map observations into phenomenon, the rules are determined by the CP, while the rules are enforced by the DP. Similarly, the utility functions, policies, or defining the parameters of the game are CP functions, while the actual decision making is a DP function. The actual invocation of the action is a DP function.

The network-level OODA loop is similar to the device-level loop, though it has a global situational awareness derived from information provided by the network elements. So the OODA description above is modified to the following:

- In the observe phase, the DP function is still the actual collection of the data. Determining which type of information to collect and how to trade off the power and energy needs of an element against that of the normal computation would be the CP functions.
- In the orient phase, the mapping of observations into phenomena is a joint process, where some local fusion and processing may occur at the DP (to minimize network loading) and then part-processed observations are sent to the SDNC that will carry out global fusion and processing (e.g., threat correlation). Again, the CP function is the definition of the parameters in the algorithms that can enable the mapping of observation to phenomena, including the division of processing responsibilities.
- In the decide phase, the policy rules provided to the SDNC via the NBI are implemented in the CP which leads to decisions covering policies/rules/configuration on traffic handling, security configuration, and measurement functionality (to support the observe phase in both inner and outer loops).
- In the act phase the policies/rules/configurations are sent via the CCI to the individual DPs.

11.4.3 Orientation Phase Algorithms

As mentioned in Sect. 11.4.2, the primary task of the orientation phase is that of mapping observed data into a higher-level concept of phenomenon. The essential algorithm in the orientation phase is to take the input data and map it into a phenomenon. The input data consists of the network traffic (e.g., packet header logs, packet payload information), system logs (indicating errors and alerts), management data (e.g., information collected from SNMP MIBs or other management information on a device), etc. The phenomenon is a high-level description of the security situation in the environment, e.g., determining if there is an intruder node, if there is a misconfiguration in the environment, if a network or link is being overwhelmed because of a denial of service attack, etc.

The orientation phase algorithm can be viewed as a classification problem where the input data needs to be mapped into one or more classes, each class indicating the existence of one particular phenomenon. The classification problem can be solved by a variety of approaches, many of which borrow heavily from the field of artificial intelligence. In the classification problem, the input data is distilled down into one or more sets of features, and a combination of those features can be used to determine which phenomenon is being experienced. As an example, the input traffic data can be mapped into features such as the values defined for configuration parameters, values measured from system and network load, and distribution parameters of network traffic (e.g., the most popular and the least popular addresses and ports used by a given machine in the network and the relative ratios among them). Each of these features can be viewed as one dimension in a multidimensional plane, with the specific value of the features determining a point in the plane. Each observation (e.g., observing the network traffic at some interval of time) provides such a point, and one needs to determine which point belongs to which class (this indicating its corresponding phenomenon).

There are a variety of algorithms that can be used to address the classification problem. These include the k-nearest neighbor algorithm, support vector machines, neural networks, decision trees, and Bayes classifiers. A survey of these algorithms and how they are used for intrusion detection are found in [10].

Another set of algorithms that can be used to map input data to phenomenon consist of root cause analysis algorithms used in system management. These algorithms map input data into symptoms (which are essentially same as the features described earlier for AI-based algorithms) and try to figure out the root cause (or the phenomenon) which is causing the symptoms to exhibit themselves. These algorithms include network topology analysis, rule-based methods, decision trees, dependency graphs, and case-based reasoning. These algorithms and their use for root cause analysis are described in more detail in [11].

11.4.4 SDN for OODA-Based Cybersecurity

In the SDN architecture for cybersecurity, cybersecurity software on various devices implements the DP functions. These use the SBI to get their configuration and policies from the SDN agent on the same device. The SDN agent uses the CCI for its configuration, rules, and policies from an SDNC, which provides them with the right information needed for the DP operation. Each device implements its device-level OODA loop and gets its CP information using the CCI. Similarly, the network-level OODA loop also needs to get its CP information, which is also provided by the SDNC using the CCI. The contents of the CP information would be very different for the network-level OODA loop and the device-level OODA loop, but they can use the same protocol to communicate with the SDNC. The structure is as shown in Fig. 11.7.

The SDN approach ensures that the rules and configurations determined by the SDNC are provided to different elements. In tactical environments, where bandwidth is limited, disruption tolerant approaches would be needed to keep the configuration parameters and policies of different elements in synchronization with the values determined by the SDNC.

In the two-tier architecture, the DP consists of both the processing in the device-level OODA loops at various devices and the network-level OODA loop. The device-level OODA loops can communicate with the network-level OODA loop using data flows that do not go through the SDNC1. These data flows are not shown in Fig. 11.7.

As an example of data path communication between the various device-level OODA loops and network-level OODA loop, consider an environment where one wants to perform an intrusion detection function for the network, using an approach like deep packet inspection. In this approach, protocol headers at higher levels are reconstructed, and they require that network exchanges, both from a client to a server (forward path) and from a server to a client (reverse path), be observed and

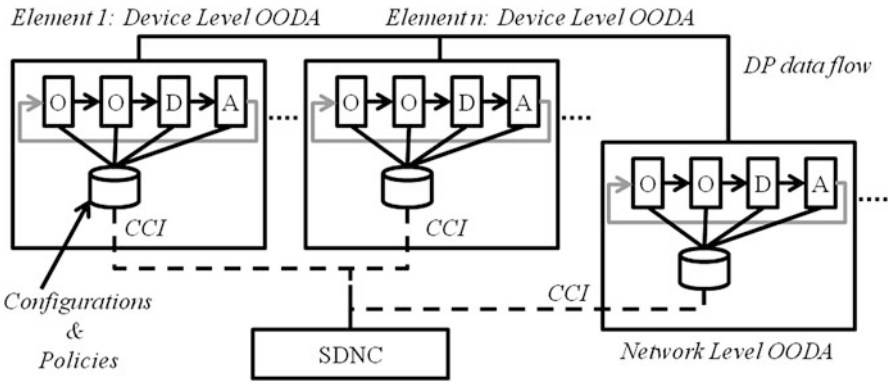

Fig. 11.7 OODA loop-based cybersecurity architecture using SDN

the results be combined to infer the progress of the higher-level protocol. However, in Internet Protocol (IP)-based communication, it is not unusual to have the forward path of a network exchange using the Transmission Control Protocol (TCP) to go through a set of devices that is different in the reverse path. In those cases, a device can only see half of the total packets being exchanged. The device-level OODA loop seeing the packets of the forward path and the device-level OODA loop seeing the packets on the reverse path can send a copy of the packets to the network-level OODA loop, which now has the information on both paths to perform the complete intrusion detection function processing.

Note that both the device-level and network-level OODA loops are data plane functions. The network-level OODA loop may or may not be collocated with the SDNC. When it is collocated with the SDNC, the benefit is that the number of points of vulnerability which can be used to attack the situational awareness system is reduced. When it is located on a separate device, the SDNC needs to be aware of that location and provide the appropriate CP information to that location. The advantage of separating the two is that the network-level OODA loop, which requires more resources, can be managed for scalability. An example of such manageability would be the ability to create multiple processes or virtual machines that perform the network-level OODA loop and to adjust the number of such processes and virtual machines depending on the amount of work needed. For the control path functions of SDNC, such scaling up and down for performance is not likely to be needed. The choice between the two modes, collocating the network-level OODA loop and the SDNC or having them on different machines, is dependent on the environment.

11.5 Coalition Tactical Environments

In a coalition environment, networks from two or more organizations need to work together. In a military coalition, the militaries from two or more nations need to come together to perform a joint mission [5]. In civilian coalitions, two independent agencies, e.g., firefighters and policemen need to work together.

The current state of the art is to have such collaboration mostly in the backend or base environments. Using SDN and the new architecture we propose, we can enable collaboration among coalition partners in the tactical environments as well. In general, coalition operations would set up their environments independently and have some level of network connectivity among them. They may have one or more tactical environments within each nation's network. In a typical coalition operation, a community of interest (CoI) is dynamically formed to conduct joint coalition operations. In a military context, the CoI can be an ad hoc team consisting of several coalition partners executing many concurrent missions including border/perimeter reconnaissance and surveillance, camp site surveillance, and detection/classification of human activities in concealed/confined spaces or locations of human infrastructures. In a civilian context, a CoI may be formed to search for missing people, rescue people from a derailed locomotive, or handle a fire in a high-rise building.

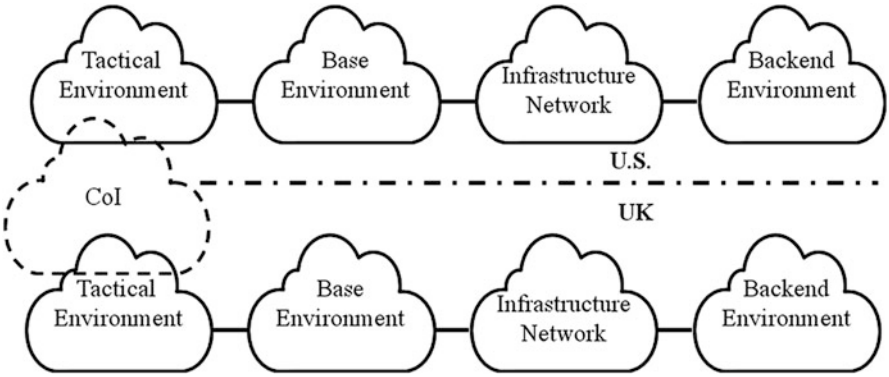

Fig. 11.8 Simplified model of a coalition network between the USA and the UK

A CoI brings together a set of assets, specific missions, and sets of policies that govern information security and sharing of information. The CoI environment would be built by combining assets from the tactical environments of multiple coalition partners, i.e., the dynamic CoI would take some assets from all of its partners in order to conduct its mission. One such sharing arrangement is shown in Fig. 11.8 where a dynamic CoI is formed between a US and UK coalition, e.g., when a joint patrol is formed to conduct surveillance in a specific area. In other cases, the CoI may also share assets from the base and other environments, including access to the back end.

When such a dynamic CoI is formed, assets from different partners may be shared. Each of the two nations may have policies limiting how the assets are shared, as well as how information from an asset may be shared with coalition partners.

11.6 Alternative Coalition SDN Architectures

In a coalition environment, we need to have a solution which brings together SDNC belonging to many different partners and have the resulting system work together seamlessly. Each partner in a coalition is likely to have a SDNC it operates and controls. In this section, we look at the various alternative options that can be used to coordinate different SDNC belonging to different partners.

For ease of notation, we are describing the coalition architecture as if there are two partners, the USA and the UK which are making a dynamic CoI. However, the architectures that are described here are applicable for a coalition of multiple partners (more than three) and may not include either of the two named countries. The two countries are just used as a short convenience for two coalition members.

Figure 11.9 shows one possible approach to support dynamic CoI in coalition SDN networks, which is to define a dynamic SDNC that is designated specifically for the CoI being supported. In this approach, if the USA and the UK need to form a

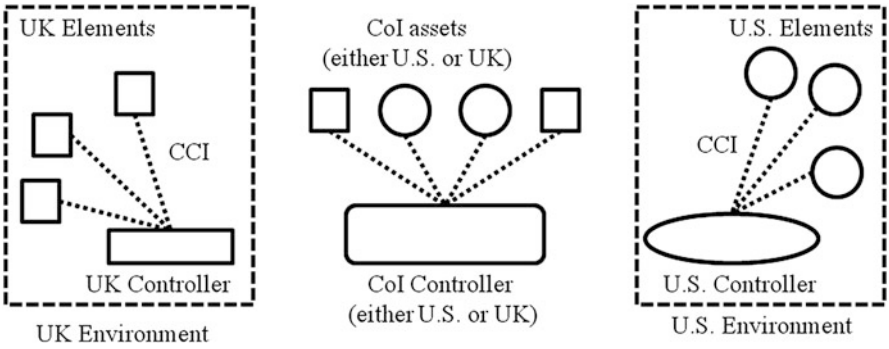

Fig. 11.9 Simplifying coalition CoI to a single organization CoI

tactical network with their assets, they designate one of their SDN controllers as the one to be used for the CoI. The different assets belonging to the coalition partners get configured and enabled by the CoI SDNC, as if they were part of the tactical environment for a single organization, as described in Sect. 11.2.2.

The advantage of this approach is that the operational logic and mechanics of the CoI is no different than that of the single organization network. The disadvantage of the approach is that it requires all nodes to interoperate with the SDNC. If the nodes from both countries use the same protocol, it is a nonissue. However, if the assets from different countries do not have the same protocol for communicating with the controller, the only viable solution is to use assets from only one country. Thus, the main decision in forming a CoI becomes which country/organization should be the one providing the assets. As a result, this approach does not enable efficient sharing of resources.

An alternative approach uses multi-domain multi-broker architecture in which the SDNCs retain their autonomy and communicate via a broker layer [6]. In such an architecture, an additional broker acts as the mechanism for enabling the decision making for SDNC from different countries, as shown in Fig. 11.10. In addition to the standard SDN controllers, a broker is introduced which acts as another layer providing the top-level hierarchy for coordinating SDN brokers. The layout of the broker and its relation to the SDNC of different coalition partners is illustrated in Fig. 11.10.

The advantage of this architecture is that each asset talks to the controller of their own organization, eliminating the challenges associated with interoperability. The broker provides the ability for the SDNC of each organization to work with each other, in effect, becoming a super controller. The main challenge with this approach is the issue associated with the operation of the broker. The coalition member operating the broker has a significant advantage in controlling the CoI compared to other partners. The issue of deciding which partner ought to run the broker can easily become very contentious.

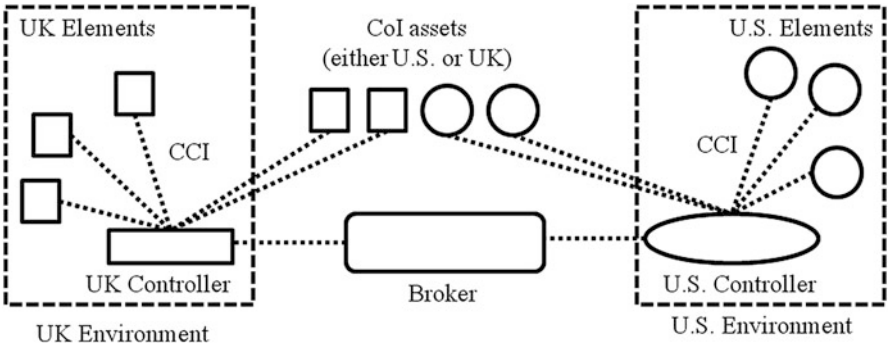

Fig. 11.10 Broker architecture for coalition SDN controllers

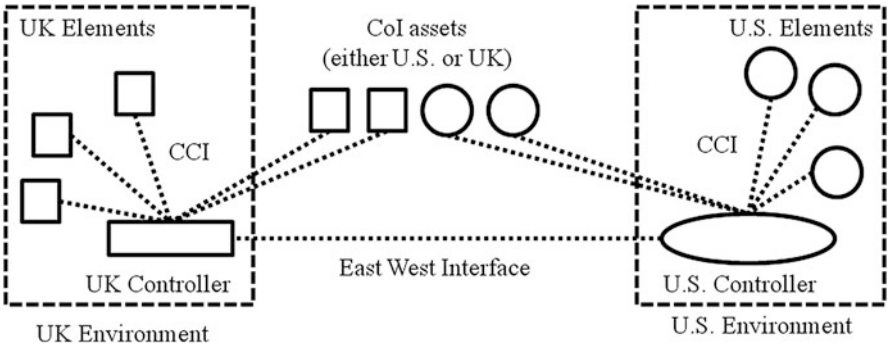

Fig. 11.11 Federated architecture for coalition SDN controllers

Another approach creates a distributed environment in which the broker is eliminated, while the equivalent functionality is provided by the collection of each country SDNC. This approach, illustrated in Fig. 11.11, avoids the tricky issue of control over the broker. In this approach, when a device requires direction for its DP actions, the controllers negotiate between themselves on an appropriate response that is forwarded to the device by the owning nation's SDNC. The distributed approach requires an east-west interface connecting different controllers and is operationally more secure since no additional elements are introduced which can act as a point of vulnerability.

These three solutions present alternative approaches for handling the issue of federation across different coalition controllers. The choice of the right solution depends on the level of trust among different partners and the degree of standardization between the nodes and SDNC. When SDNC and the nodes use the same interface, reducing the problem to a single organization system for the CoI would work. In other cases, the choice depends on the level of trust among coalition partners. When one partner is trusted to operate a broker, the broker-based approach will be most appropriate. When partners only trust each other partially, the distributed east-west approach is more suitable.

A comparison of the three different architectures is provided in [9]. The analysis performed there shows that the interoperability among different coalition environments is enhanced significantly by the federated and brokered architectures, as compared to the approach of simplifying the problem to a single organization architecture. Furthermore, from a complexity perspective, the simplification approach is the one with least amount of complexity. The broker approach is more complex, and the federated architecture is the most complex solution among the three. From a trust relationship perspective, the simplification approach and the broker approach require more trust in a single organization than the federated approach. From a standardization perspective, the simplification approach requires a higher degree of standards to be defined than the brokered or federated architectures.

11.6.1 Federated SDN-Based Cybersecurity for Coalitions

In a coalition environment, partners do not fully trust each other. As a result, the federated architecture described in Fig. 11.11 is the preferred solution for many coalition tactical environments. In these environments, each country network is likely to have their own SDNC, and all of the SDNCs need to be federated together to create a completely functional system for the overall network. As mentioned previously, that implies that the SDNC needs to be augmented with not just a north-south interface between the elements and the SDNC in individual country networks but with an east-west interface that is used to exchange information between the individual SDNCs. In this respect, the architecture we propose is similar in principle to coalition operations for ISR assets [7] and federation of military networks [8].

One way to define the east-west interface is to use a mechanism based on distributed systems such as Hyperledger [12]. Hyperledger is a system which allows tracking of transactions among different parties which all maintain a peer-to-peer relationship with each other and implements a distributed consensus protocol for all peers to determine whether or not a transaction has happened. An architecture for software defined coalitions based on Hyperledger is described in [13] and provides one of the ways in which the federated architecture described here can be implemented.

The architecture of a system with controllers from both the USA and the UK is shown in Fig. 11.12. In the figure, the oval and circular boxes represent assets belonging to the USA, while the square and rectangular boxes represent assets belonging to the UK. Each asset runs a device-level OODA loop, and let us assume that the network-level OODA loop is collocated with the controller in both countries. However, as mentioned earlier, such collocation is not strictly necessary.

The controllers in each of the individual networks are responsible for providing the policies, configurations, and parameters that drive the operation of each of their elements. The OODA loop implemented within the US elements and the UK elements could be quite different, with the use of different approaches in each of the individual country elements.

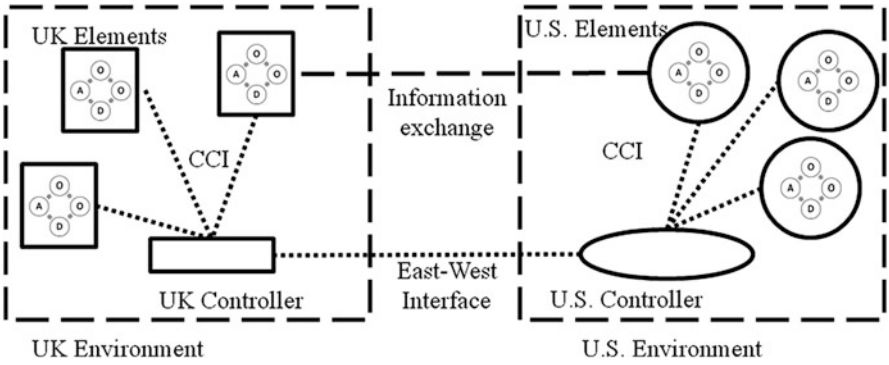

Fig. 11.12 OODA loop-based cybersecurity architecture using SDN for coalition networks

Devices in each of the two countries used the CCI interface to talk to their respective controllers. While the USA and the UK are not obligated to use the same protocol, a common protocol such as one based on a REST interface to harmonize policies and parameters of different elements is likely to be used in each nation. Nevertheless, the choice of specific names of variables and parameters, as well as policy format and specifications, are likely to be different in each nation. The east-west interface provides a mechanism for the controllers to work and interoperate with each other and to set up managed information exchange points between partner networks. This interface can be used to share policies or negotiate dynamic policies when CoI are formed dynamically.

Note that the east-west interface is used for control plane functions and to manage the data plane connections. In a coalition network, there may be multiple interconnections between the nodes for actual data exchange. Direct links may be established between US and UK nodes, if allowed by the applicable policies. The thin dashed line marked information exchange between the two nodes in Fig. 11.12 shows one such possible data path. Several of these data paths can be used in a coalition environment. As an example, if a UK node happens to be closer to several US nodes, it may choose to route packets using one or more of those US nodes instead of trying to connect only to the UK nodes. The control plane interconnection and data plane interconnections can be very different in these cases. This scenario is shown pictorially in Fig. 11.13, where the solid lines indicate the data flow used between the US and UK elements to implement an efficient routing mechanism (assuming that the same routing protocols are supported by both nations) and the dashed lines show the control flows, where each of the elements talks to the controller of their respective countries. Data flow may happen directly between the assets, but the control path information is provided by the controller of each of the two countries to their assets.

The federated coalition SDN architecture can be used to coordinate the security threat assessment and facilitate the sharing of information among different coalition partners. The information sharing can occur among the controllers (SDNC) of the

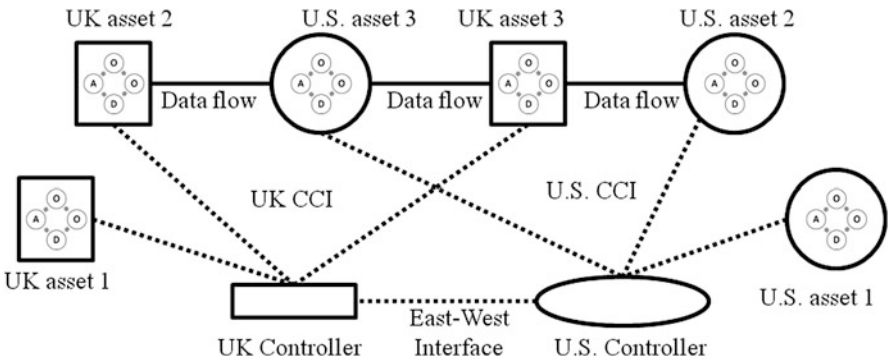

Fig. 11.13 Data path and control path flows in a coalition routing scenario

two countries, as well as directly between different nodes in the two countries. In the latter case, the SDNC would determine the policies that govern these direct exchanges but do not necessarily be on the path of actual data exchange.

As a very simple example, let us consider the case where a rogue terrorist is trying to launch an attack on the UAVs that are operated by coalition partners in a theatre of operation. Let us also assume that the terrorist has been able to determine the frequency at which commands are issued to the UAV and is trying to launch a scanning attack to determine if any communication port in the UAV is vulnerable. The USA may have detected the terrorist probes, and the US controller has installed a rule for orientation that maps more than three probes on illegal ports from a device to mark that device as unauthorized entity to be added to a blacklist. The UK detection module, however, may have ended up with a policy that locates the spatial region of the terrorist and in those regions disable all external communication and operate using a disconnected operation mode.

When a dynamic CoI is formed in which the USA and the UK both contribute UAVs for the operations, the controllers for both nations can share the policies they have formed with one another. This enables the UAVs for the CoI, which may have come from either country, to install the security policies which enable the joint insights from both nations to be used. The US UAVs can get insights about the vulnerability region in the theatre, while the UK UAVs get additional rules to learn the address of the device and block them dynamically even when exposed outside that region.

11.7 Conclusions

Coalition tactical networks are composed of different networks of two or more nations coming together for securing a mission in tactical arena. They may use heterogeneous networks using mixes of (1) handheld units, ISR devices,UAVs, (2)

fixed infrastructure wireless (cellular), (3) infrastructure-less wireless (MANET), (4) satellites, and (5) private access points. In addition they may also leverage commercial networks in the area.

SDN works on the principle of separating control plane from data processing operations and is commonly implemented using a central controller, which provides guidance to individual network elements on how they ought to execute their data processing operations. In a coalition setting, two SDN controllers working across two networks act to control and coordinate operations of their data processing, while maintaining control over their individual networks.

In this chapter, we have discussed how we can utilize the principles of SDN to improve cyber situational awareness in coalition environments. In various military networks, the task of determining situational awareness in represented as an implementation of the OODA (observe-orient-decide-act) loop. In the context of security using SDN principles, we can draw an analogue of data plane and control plane in the context of cyber situational awareness. The data plane for situational awareness can be defined as comprising the set of elements that implement the actual OODA loop. The control plane for situational awareness can be defined as comprising the set of elements that provide the configuration, background knowledge, and configuration required by the data elements.

In this chapter, we have (a) introduced tactical coalition networks, (b) presented an architecture for applying SDN principles to address the task of cyber situational awareness for network security, (c) illustrated how the architecture can be used to understand the current situation for a cybersecurity threat, and (d) discussed alternative architectures for cooperation between different SDN controllers belonging to various coalition partners.

Review Questions

- What is a tactical environment?
- What are the key differences between a tactical environment and a backend environment like a data center?
- What are the different types of segments that make up a single organization network?
- What are the drawbacks of using the central controller approach for tactical environments?
- What are the different life-cycle stages for a node in a tactical environment?
- What unique issues are introduced by coalition networks in a tactical environment?
- What are the merits and demerits of collocating the network-level OODA loop data path functions and SDNC?
- What are the disadvantages in a coalition tactical network if all information exchange between partner environments is forced to go through their respective SDNC?

- Compare the benefits and drawbacks of the three approaches for coordination between the controllers of coalition partners?
- What are the typical control plane and data plane operations for a tactical environment node during its operational stage?

Acknowledgments This research was partially sponsored by the US Army Research Laboratory and the UK Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defence, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation hereon.

References

1. McKeown N et al (2008) OpenFlow: enabling innovation in campus networks. ACM SIGCOMM Comput Commun Rev 38(2):69–74
2. Willig A, Kubisch M, Hoene C, Wolisz A (2002) Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer. IEEE Trans Ind Electron 49(6):1265–1282
3. Brehmer B (2005) The dynamic OODA loop: amalgamating boyd's OODA loop and the cybernetic approach to command and control. In: Proceedings of the 10th international command and control research technology symposium
4. Ye F et al (2012) MECA: Mobile Edge Capture and Analysis middleware for social sensing applications. In: Proceedings of the 21st International Conference on World Wide Web. ACM
5. Verma D (2010) (ed) Network science for military coalition operations: Information Exchange and Interaction. IGI Global
6. Castro A et al (2016) Brokered orchestration for end-to-end service provisioning across heterogeneous multi-operator (multi-AS) optical networks. J Lightwave Technol 34(23)
7. Calo S et al (2009) Technologies for federation and interoperation of coalition networks. Information Fusion, 2009. FUSION'09. 12th international conference on. IEEE
8. Sørensen E (2014) SDN used for policy enforcement in a federated military network
9. Mishra V, Verma D, Williams C Marcus K (May 15–16, 2017) Comparing software defined architectures for coalition operations. In: Proceedings of the international conference on military communications and information systems, Oulu, Finland
10. Tsai C, Hsu Y, Lin C, Lin W (2009) Intrusion detection by machine learning: a review. Expert Syst Appl 36(10):11994–12000
11. Verma DC (2009) Principles of computer systems and network management. Springer, Heidelberg, Chapter 6, pp 143–155
12. Cachin C (2016) Architecture of the Hyperledger blockchain fabric. In: Proceedings of the workshop on distributed cryptocurrencies and consensus ledgers
13. Verma D, Desai N, Preece A, Taylor I A block chain based architecture for asset management in coalition operations. In: Proceedings of SPIE Defense + Commercial Sensing Symposium, Anaheim, CA, April 2017

Vinod K. Mishra received his PhD in Physics from the State University of New York (SUNY) at Stony Brook in 1983, with area of focus in theoretical nuclear physics. Earlier he got his Master of Science from Indian Institute of Technology, Kanpur (1977), and Bachelor of Science (Physics Honors) from Science College, Patna (1975). He was a postdoctoral researcher at various universities and research institutions before joining Lucent Technology Bell Labs, where he

worked in many areas of optical and wireless networking. Later he came to Defense Information Systems Agency (DISA) and focused on advanced networking technologies. Currently he is a team leader at US Army Research Laboratory (ARL) conducting research in software-defined networking, dynamic optical networking, and quantum communication. He had published a book entitled *An Introduction to Quantum Communication* in 2016.

Dinesh C. Verma is an IBM fellow and manager of the Distributed Cognitive Systems area at IBM TJ Watson Research Center, Yorktown Heights, New York. In this role, he leads a research team creating new technologies that intersect the domain of cognitive computing, Internet of Things, and distributed systems and networks. He received his doctorate in Computer Science from University of California at Berkeley in 1992, Bachelors' in Computer Science from Indian Institute of Technology, Kanpur, India, in 1987, and Masters in Management of Technology from Polytechnic University, Brooklyn, NY, in 1998. He holds over 100 US patents and has authored over 100 papers and ten books in computer science. He is a fellow of the IEEE and a fellow of the UK Royal Academy of Engineering and has served in various program committees, IEEE technical committees, editorial boards, and managed international multi-institutional government programs.

Christopher Williams graduated from Oxford University with a first in engineering science and subsequently gained his PhD from Bristol University on the topic of chaotic waveforms for communications. Alongside periods in industry (research manager for Fujitsu) and academia (research fellow at Bristol University), much of his career has been in government defense research (Dstl and predecessors). His areas of expertise include novel waveforms, communications signal processing, modulation and coding, cognitive radio, and dynamic spectrum access.