ARL-TR-8842 ● Oct 2019

An Overview of ARL's Sensor Information Testbed Collaborative Research Environment (SITCORE) Concept of Operations and Environment

by Kelly Bennett and Joe Ryder

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-8842 ● OCT 2019

An Overview of ARL's Sensor Information Testbed Collaborative Research Environment (SITCORE) Concept of Operations and Environment

Kelly Bennett
Sensors and Electron Devices Directorate, CCDC Army Research Laboratory

Joe Ryder
CACI International Inc, Arlington, VA

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
October 2019	Technical Report	September 2018–September 2019

4. TITLE AND SUBTITLE
An Overview of ARL's Sensor Information Testbed Collaborative Research Environment (SITCORE) Concept of Operations and Environment

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Kelly Bennett and Joe Ryder

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
CCDC Army Research Laboratory
ATTN: FCDD-RLS-SA
2800 Powder Mill Road, Adelphi, MD 20783-1138

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-TR-8842

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The US Army Combat Capabilities Development Command Army Research Laboratory (ARL) Sensor Information Testbed Collaborative Research Environment (SITCORE) is an initiative of the Information Science campaign within ARL and has the potential to support the ARL Open Campus effort and other collaborative research as well as other efforts, including artificial-intelligence and machine-learning research. Along with the Automated Online Data Repository, these initiatives create a research laboratory and testbed environment focused on sensor data and information fusion. This report serves as the technical concepts and operations overview for the SITCORE hosted on Amazon Web Services for the ARL.

15. SUBJECT TERMS
sensor testbed, information fusion, research collaboration, System Security Plan, SSP, authority to operate, ATO, Sensor Information Testbed Collaborative Research Environment, SITCORE, SITCORE Technical Plan to Operate, Automated Online Data Repository, AODR, Amazon Web Service, cloud computing

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Kelly Bennett
Unclassified	Unclassified	Unclassified	UU	31	19b. TELEPHONE NUMBER (Include area code) (301) 394-2449

Contents

List of Figures

List of Tables

1. Introduction

The US Army Combat Capabilities Development Command Army Research Laboratory (ARL) Signal & Image Processing Division, within the Sensors and Electron Devices Directorate, collaborates with US and coalition partners on numerous sensor-data-research initiatives. ARL has established an Open Campus concept to facilitate wider collaboration and successfully share and transition research data throughout the larger research community. ARL is searching for solutions to facilitate joint research, software development, and a collaborative testing environment in conjunction with the ARL Open Campus initiative.[1]

The ARL Sensor Information Testbed Collaborative Research Environment (SITCORE) is an information science campaign initiative within the Open Campus concept. SITCORE's primary objective is to establish a collaborative research laboratory and testbed environment focused on sensor data and information fusion. For the ARL Signal & Image Processing Division, this objective translates to making sensor data available to the sensor research community to develop algorithms based on a Common Data Representation and use data to test within the community.

SITCORE's goal is to create a collaborative environment for ARL and external partners. These external partners can include postdoctoral students, NATO and coalition partners, industry, and academia. SITCORE provides an environment for dispersed and diverse researchers and developers to collaborate, share, and benefit from each other's work. The intention is to enable "end-to-end" algorithm development and experiments to validate and demonstrate results. SITCORE will use an Open Campus environment; therefore, the system is intended for non-sensitive information.

2. Scope

2.1 Information Sensitivity

Information processed or generated in SITCORE does not include personally identifiable or classified information. The highest classification of data processed or generated in SITCORE is publicly releasable.

2.2 System Security

System security policy shall be in accordance with the Risk Management Framework for Department of Defense (DOD) Information Technology (IT) directive.[2]

2.3 Technical Requirements

SITCORE is hosted on the Amazon Web Services (AWS) cloud infrastructure. A number of AWS services are utilized to host SITCORE. The overarching SITCORE Uniform Resource Locator (URL) domain is www.sitcore.net; each tool making up the SITCORE suite is described in the following sections.

2.4 Amazon Web Services (AWS)

According to Amazon:

> Amazon Web Services is a secure cloud services platform, offering compute power, database storage, content delivery and other functionality to help businesses scale and grow … with increased flexibility, scalability and reliability.

2.5 Amazon Elastic Compute Cloud (EC2)

Amazon EC2 servers provide scalable computing capacity in the AWS cloud. SITCORE currently comprises 12 different EC2 instances as shown in Table 1.

Table 1 SITCORE EC2 instances

Instance name	Description	Production cycle phase	Time running (estimated %)	Volumes encrypted at rest?	Associated RDS instance
aws-cloud9- SITCORE-Cloud9-7d961ebf5ec34b23814f8ed74403a280	Enables Cloud9 functionality	Production	5%	No	NA
MMSDB_Dev	A copy of the non-FOUO contents of MMSDB	. . .	100%	No	SITCORE-MySQL
osTicket-stand-alone	Runs osTicket: The software behind support.sitcore.net	Production	100%	Yes	SITCORE-MySQL
SITCORE	Home page: https://www.sitcore.net	Production	100%	No	NA
SITCORE-DevOps	Holdover from when SITCORE was hosted on a single EC2 instance. Runs Xwiki, Jenkins, and Redmine	Production	100%	Yes	SITCORE-PostgreSQL
SITCORE-Docker	A future replacement for SITCORE-DevOps; Runs similar tools to SITCORE-DevOps, but from within Docker containers	Development	40%	Yes	SITCORE-MySQL
SITCORE-GitLab	Runs GitLab	Production	100%	Yes	SITCORE-PostgreSQL
SITCORE-Gitlab-test	A staging/testing environment for GitLab	Development/ testing	. . .	Yes	SITCORE-PostgreSQL
SITCORE-Mattermost	Runs Mattermost, a team chatting application	Production	100%	Yes	SITCORE-PostgreSQL
SITCORE-MiniDass	Supports the Mini-Dass team	Production	100%	Yes	NA
SITCORE-OpenSim-Ubuntu	Runs OpenSim virtual reality software	Production	0%	No	sitcore-mysql-opensimulator
SITCORE-WMRD	A large EC2 instance to support the ARL WMRD team	Custom	0% (as of 3/13/2018)	No	NA

Note: RDS = Relational Database Service; NA = not applicable; MMSDB = Multimodal Signature Database; EC2 = Elastic Compute Cloud; WMRD = Weapons and Materials Research Directorate.

2.6 Load Balancing

Assuming there are multiple mirror EC2 instances in the cluster, the Elastic Load Balancer (ELB) will route tasks to their respective instances. The port-forwarding rules dictate which ports the ELB will listen on and which EC2 instance port it will forward each request to. It is easiest to think about the base case: a single EC2 instance running within an ELB-managed cluster. Figure 1 shows a typical Elastic Block Storage (EBS) configuration where an inbound request is routed to a particular EC2 instance.

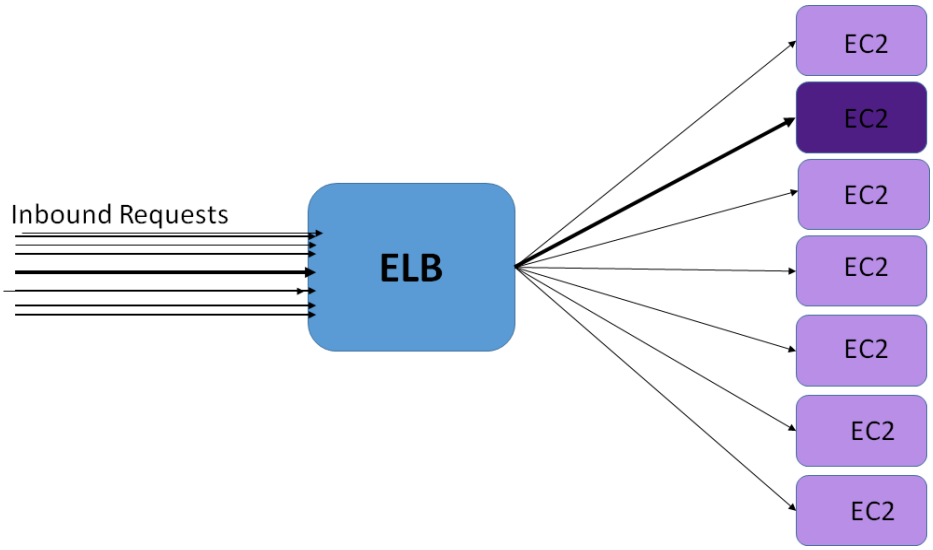

Fig. 1 ELB configuration

2.7 Application Load Balancer (ALB)

All of the load balancers within SITCORE are ALBs. Built to enable URL-based routing, their construction differs slightly from the traditional model shown in Fig. 1.

As shown in Fig. 2, ALBs route incoming application requests to different target groups based on their URL. Each target group can have one or more associated EC2 instances. Assuming all instances are healthy, the request will be routed to an associated EC2 instance using a load-balancing algorithm (such as round robin). This extension allows for autoscaling of applications should they face sudden changes in demand.

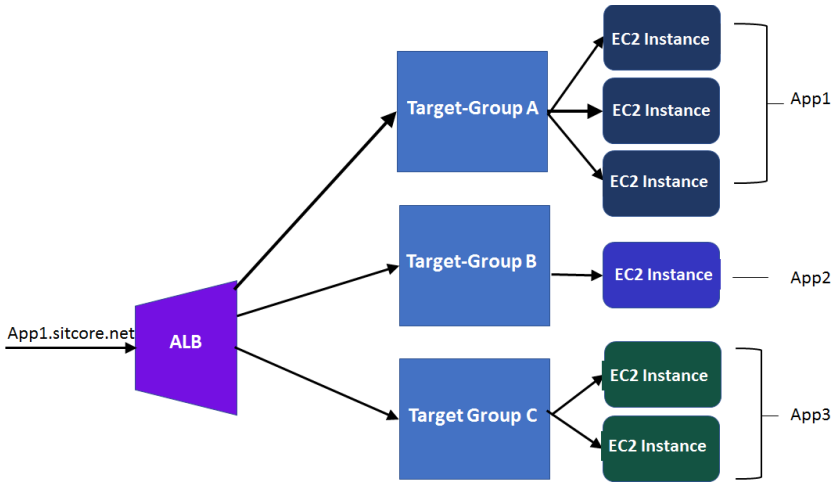

Fig. 2 ALB configuration

2.8 Load Balancers and Transit Encryption Configurations in SITCORE

To set up encryption in transit, all load balancers in SITCORE are configured to listen for secure socket layer (SSL)-encrypted communications and to perform SSL termination. This means all applications running behind a load balancer have transit encryption enabled between the load balancer and the client, but are sent over HyperText Transfer Protocol (HTTP) when forwarded from the load balancer to the server. Given that the ALB and all of SITCORE's EC2 instances (virtual servers) are enclosed within the same Virtual Private Cloud (VPC) as shown in Fig. 3, the configuration is very secure.

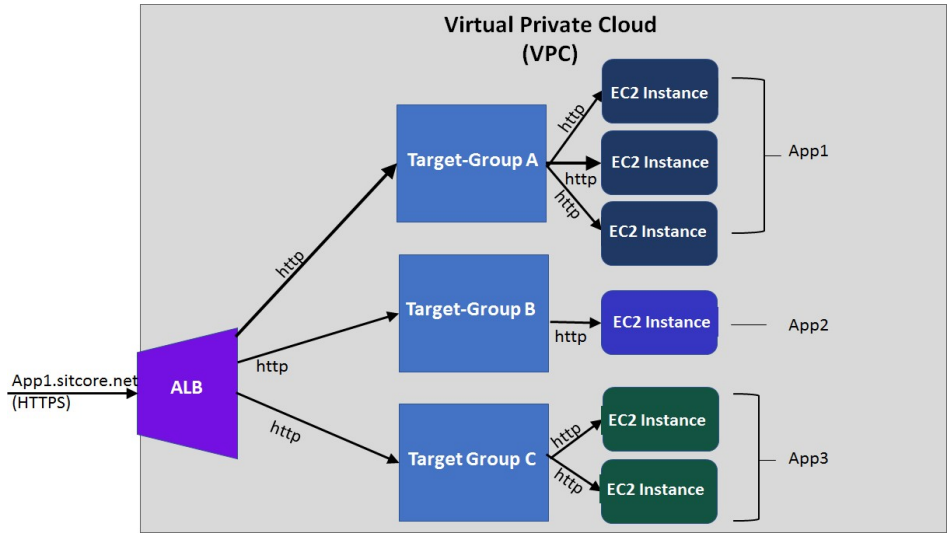

Fig. 3 ALB and EC2 instances enclosed within VPC

This security is enforced by the VPC in conjunction with the EC2 instances' configurations. The VPC blocks all inbound communications with the exception of those made to the ALBs, and each of the EC2 instances is configured to forward all external communications over HTTP to HyperText Transfer Protocol Secure (HTTPS).

Table 2 shows a listing of all of the load balancers within SITCORE along with their target groups and the status of their SSL configurations.

Table 2 Load balancers within SITCORE

Load balancer name	Target groups	SSL configured?
Extended-SITCORE	MiniDass	Yes
sitcore-home	• sitcore • DevOps	Yes
sitcore-production	• mattermost • sitcore-gitlab • sitcore-reboot	Yes
support-at-sitcore	osTicket	Yes
Test-Box	• sitcore-gitlab-test • sitcore-test	Yes

2.9 Amazon Simple Storage Service (S3)

Amazon S3 allows for secure, durable, and highly scalable object storage. Current S3 buckets are as follows:

- sitcore: general parent bucket to hold any future SITCORE data not stored on ownCloud and Phabricator, as needed; bucket currently empty
- sitcore-ownCloud: currently used for all ownCloud data storage
- sitcore-phab: meant for Phabricator, but currently not needed as all data are on Relational Database Service (RDS); bucket currently empty

2.10 Amazon RDS

Amazon RDS provides cost-efficient and resizable capacity. SITCORE has four different EC2 instances to handle all database (DB) processing as shown in Table 3.

Table 3 Amazon RDS within SITCORE

Engine	DB instance	Status	Class	VPC	Multi-AZ	Encrypted
MySQL	sitcore-mysql	Available	db.t2.medium	vpc-faef7b9e	Yes	No
MySQL	sitcore-mysql-opensimulator	Stopped	db.t2.medium	vpc-faef7b9e	No	No
Postgresql	sitcore-postgresql	Available	db.t2.medium	vpc-faef7b9e	Yes	No
Postgresql	test-postgresql	Stopped	db.t2.micro	vpc-faef7b9e	No	No

Note: DB = database; VPC = Virtual Private Cloud; Multi-AZ = multiple availability zones.

2.11 Amazon EBS

Amazon EBS provides persistent block-level storage volumes for use with Amazon EC2 instances in the AWS cloud. Each Amazon EBS volume is automatically replicated within its availability zone to protect the system from component failure. SITCORE does not currently store data in EBS volumes. Table 4 shows the EBS within SITCORE.

Table 4 EBS within SITCORE

Name	Volume ID	Size	Volume type	IOPS	Availability zone	Attachment information	Monitoring	Encrypted
sitcore-ebs	vol-5eddf8bd	50 GiB	gp2	150/3000	us-east-1b	**i-13b932a5**:/dev/sdd (attached)	Yes	Not Encrypted

Note: EBS = Elastic Block Storage; ID = identification; IOPS = input/output operations per second.

2.12 Amazon Autoscaling Groups

Autoscaling maintains application availability and allows SITCORE to scale Amazon EC2 capacity up or down automatically according to predefined conditions. It automatically deploys replicated instances of any EC2 instance whose load approaches CPU or memory limit.

2.13 Amazon ELBs

ELBs automatically distribute incoming application traffic across multiple Amazon EC2 instances in the cloud. It achieves greater levels of fault tolerance in SITCORE applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic. This service works in conjunction with autoscaling. As new instances are spun up, the load balancer starts distributing traffic to them as well.

2.14 Amazon Elastic Internet Protocol (IP)

An Elastic IP address is a static IP address designed for dynamic cloud computing. With an Elastic IP address, the failure of an instance or software can be masked by rapidly remapping the address to another instance in the SITCORE account without any loss of availability experienced by the user.

Note: Awaiting final production implementation to be configured.

2.15 Amazon Route 53

Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. Amazon Route 53 effectively connects user requests to infrastructure running in AWS—such as Amazon EC2 instances, ELBs, or S3 buckets—and can also be used to route users to infrastructure outside of AWS. The SITCORE domain is www.sitcore.net.

3. Network Topology

Figure 4 shows the network routing of the SITCORE cloud environment. The configuration takes advantage of VPC peering to connect two VPCs and route traffic between them using private networking addresses.

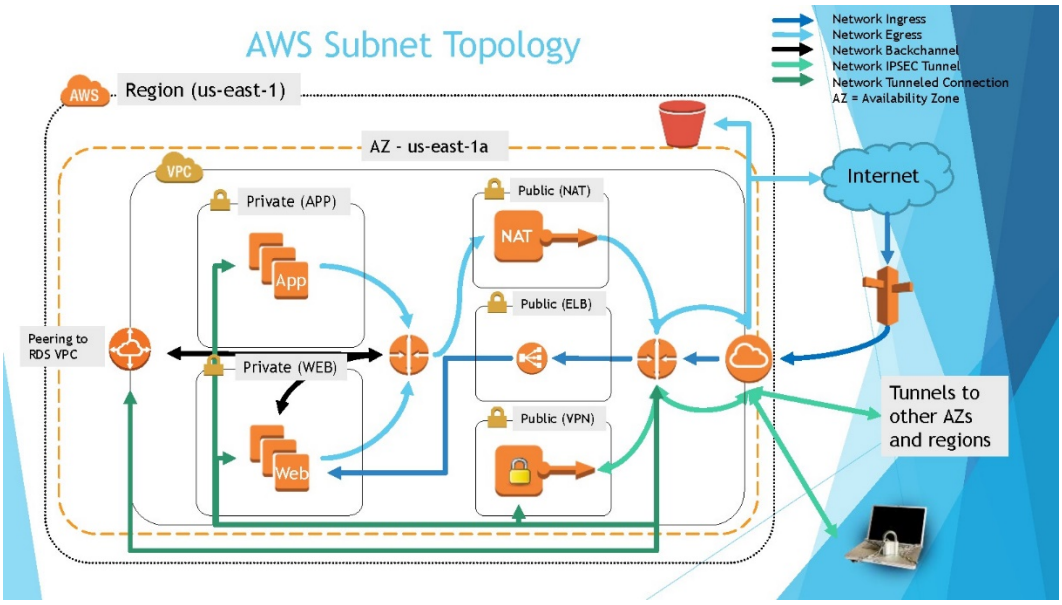

Fig. 4 AWS Subnet Topology used for the SITCORE environment showing subnets and routing of public and private instances

Figure 5 shows Route 53 and AWS EC2 used by SITCORE. Route 53 is a highly available and scalable DNS service developed by AWS. Combined with the ALB, it creates a robust environment for SITCORE EC2 instances.

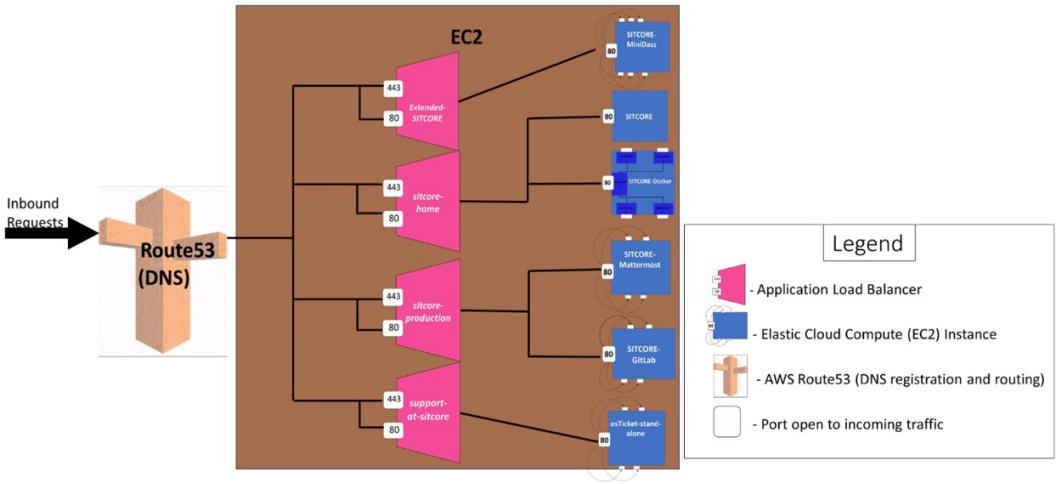

Fig. 5 Route 53 DNS, ALB, and EC2 instances used in SITCORE network topology

4. Operations

4.1 Obtaining Access

The system is a web application accessible from modern Internet browsers. Access to the application is controlled by standard Public Key Infrastructure authentication. Users register with the site by creating an account. Once an account is approved, users access the site via a DOD common access card (CAC) or similar approved non-DOD smart certificate. The application uses a role-based system to control access to the site's pages and resources. There are four roles: GUEST, USER, ADMIN, and BANNED.

4.2 Account Management

A USER has the ability to access all tools in SITCORE and the ability to upload and download GITLAB code projects in the domain in which they have been approved for access, view his/her historical requests, approve other users' requests when acting in the capacity as a government sponsor or approver, and view or update his/her user account information.

An ADMIN, in addition to having the same access as a USER, has elevated permissions within the system. An ADMIN can perform the following additional actions: approve or reject new user or organization registrations; approve or reject

requests; edit projects; delete users, organizations, or datasets; search all users, organizations and requests; add new datasets; and update site settings.

An individual who submits a user-registration request is initially assigned a role of GUEST. To approve the registration, an ADMIN will assign the individual either the USER or ADMIN role. To reject the registration, an ADMIN will assign the individual the BANNED role.

Accounts that have been inactive (not logged into) for 180 d will be automatically disabled. Inactive users will receive a notification by email 7 d prior to disabling their account, providing an opportunity to log into the system and reactivate the account. All user accounts will automatically expire after 2 years. Users will receive a notification by email 1 month prior to the expiration and will be able to request that site administrators extend the expiration date of their account. Expired, inactive, or unapproved accounts are assigned the BANNED role.

The application's main page is a webpage with links to Services and Help. From this page, users can gain access to the various services supported by SITCORE: Gilab, Redmine, Jenkins, and Mattermost. Other pages accessible via the webpage are Help: XWIKI Forum, Support, and Contact Form.

AWS Identify and Access Management (IAM) securely controls access to AWS services and resources for users. IAM enables SITCORE administrators to create and manage AWS users and groups, using permissions to allow and deny access to AWS resources. By default, IAM users do not have permission to create or modify Amazon Elastic Container Service (ECS) resources or perform tasks using the Amazon ECS application program interfaces (API). This means they also cannot do so using the Amazon ECS console or the AWS Command Line Interface (CLI). To allow IAM users to create or modify resources and perform tasks, you must create IAM policies that grant IAM users permission to use the specific resources and API actions they will need, and then attach those policies to the IAM users or groups that require those permissions. When you attach a policy to a user or group of users, it allows or denies the users permission to perform the specified tasks on the specified resources. Likewise, Amazon ECS container instances make calls to the Amazon ECS and Amazon EC2 APIs on your behalf, so they need to authenticate with your credentials. This authentication is accomplished by creating an IAM role for your container instances and associating that role with your container instances when you launch them. With your Amazon ECS services, calls to the Amazon EC2 and ELB APIs are made on your behalf to register and deregister container instances with your load balancers.

An IAM policy is a JSON document that consists of one or more statements. Each statement is structured as follows:

```
{
"Statement":[{
"Effect":"effect",
"Action":"action",
"Resource":"arn", "Condition":{
"condition":{ "key":"value" } } }
]
}
```

There are various elements that make up a statement:

1) Effect: The effect can be Allow or Deny. By default, IAM users do not have permission to use resources and API actions, so all requests are denied. An explicit Allow overrides the default. An explicit Deny overrides any Allows.

   Action: The action is the specific API action for which you are granting or denying permission.

2) Resource: The resource that is affected by the action. Some Amazon ECS API actions allow you to include specific resources in your policy that can be created or modified by the action. To specify a resource in the statement, you need to use its Amazon Resource

3) Amazon Resource Name (ARN): If the API action does not support ARNs, use the * wildcard to specify that all resources can be affected by the action.

   Condition: Conditions are optional. They can control when your policy will be in effect.

4.3 Security

To enhance the security of SITCORE, two-factor authentication will be required when accessing all segments of SITCORE. CAC-enabled users will authenticate via Enterprise Access Management Service-Army (EAMS-A), which in turn will connect to AWS multifactor and AWS single sign on (SSO) to access SITCORE and will have access to government services such as the Automated Online Data Repository.[3,4] Non-CAC users will authenticate through AWS multifactor and AWS SSO and will only have access to SITCORE publicly releasable systems. Figure 6 shows a diagram of the SSO and authentication design using both two-factor (CAC/External Certification Authority [ECA]) or non-CAC users.

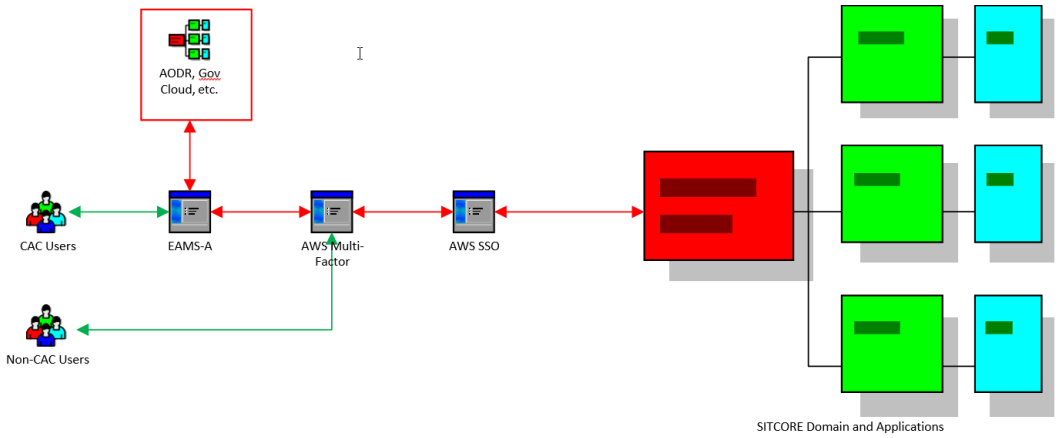

Fig. 6 SITCORE design of SSO using two-factor authentication (CAC or ECA)

4.4 Secure Socket Layer (SSL) Certificate

The SITCORE site uses an OpenSSL-generated self-signed certificate. Please refer to Fig. 2 for more information.

4.5 Public/Private Secure Shell (SSH) Key

Phabricator users generate an SSH key pair with an Rivest–Shamir–Adleman (RSA) private key on their hard drive and an RSA public key loaded onto Phabricator. AWS uses 2048-bit SSH-2 RSA keys. Jenkins authenticates with Phabricator via a Jenkins bot user and an SSH key. No other SITCORE tools use SSH keys.

4.6 AWS Security Features

AWS is a Federal Risk and Authorization Management Program (FedRAMP)-compliant Cloud Service Provider (CSP). FedRAMP is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.

AWS has completed the testing performed by a FedRAMP-accredited Third-Party Assessment Organization and has been granted two Agency Authority to Operates (ATOs) by the US Department of Health and Human Services after demonstrating compliance with FedRAMP requirements.

AWS's compliance with FedRAMP requirements was achieved based on testing performed against the stringent set of FedRAMP requirements,[5] plus additional FedRAMP security controls. AWS has been assessed and approved as a CSP for security impact Level 2: public unclassified data.

- At Level 2, all AWS US-based regions—US East/West and AWS GovCloud (US)—have been assessed by the Defense Information Systems Agency and issued two provisional authorizations after demonstrating compliance with DOD requirements. AWS's compliance with DOD requirements was achieved by leveraging the existing FedRAMP Agency ATOs. The provisional authorizations allow DOD entities to evaluate AWS's security and the opportunity to store, process, and maintain a diverse array of DOD data within the AWS cloud.

In a standard AWS cloud environment, the security responsibilities are shared between the CSP and the cloud customer. The level of CSP and customer responsibilities in this shared responsibility model depends on the cloud deployment model. SITCORE uses an Infrastructure as a Service with shared-responsibility model as follows:

- AWS responsibility: AWS operates, manages, and controls the infrastructure components from the host operating system (OS) and virtualization layer down to the physical security of the facilities in which the service operates.
- SITCORE responsibility: Customers/partners assume responsibility and management of the guest OS (including updates and security patches), other associated application software, as well as the configuration of the AWS-provided security group firewalls and other security, change management, and logging features. Security in the cloud is based on a shared security model.[6]

4.7 Firewall Configuration

Firewalls are located in front of each SITCORE server as detailed in Fig. 1. Amazon EC2 provides a firewall solution; these mandatory inbound firewalls are configured in a default deny-all mode. Only the ports necessary for the SITCORE and Open Simulator to operate have been opened.

4.8 Resilience Measures

SITCORE uses ELB and autoscaling instances to ensure redundancy and availability. ELB automatically scales its request handling capacity to meet the demands of application traffic.

ELB integrates with autoscaling to ensure there is enough back-end capacity to meet varying levels of traffic levels without needing manual intervention.

Another AWS service that SITCORE uses is Amazon S3. Amazon S3 storage provides the highest level of data durability and availability in the AWS platform. Error correction is built in, and there are no single points of failure. Amazon S3 is designed for better than 99% durability per object and better than 99% availability over a 1-year period; therefore, there is very low risk of informational loss or the ability to access AWS S3 data.

4.9 Backup Policy

SITCORE takes advantage of several of AWS's storage options. Amazon S3 provides a simple web services' interface that can be used to store and retrieve any amount of data, at any time, from within Amazon EC2 or from anywhere on the web. One can write, read, and delete objects containing from 1 byte to 5 terabytes of data each. The number of objects one can store in an Amazon S3 bucket is virtually unlimited. Data will reside on Amazon S3 in the short term (~30 d). SITCORE is configured to automatically back up snapshots of Amazon RDS DBs. Amazon stores all of these backups on S3.

However, AWS does not back up EC2 and EBS by default. These services must be backed up manually. An automated script was created for SITCORE that runs once a week taking snapshots of the EC2 instance, which contains EBS. Six weekly backups are stored, with the backup from 7 weeks ago being deleted.

After 30 d, data can be set up to be transferred to Amazon's Glacier service, which is a long-term data-storage solution. Data reside on Amazon Glacier until removal. Amazon Glacier is designed for the same durability as Amazon S3. Data are stored in Amazon Glacier as archives. An archive can represent a single file, or it could be several files. Archives are organized into vaults. Access to these vaults are controlled using the AWS IAM service.

4.10 Container Security

Running Docker containers within the execution environment of customer-controlled EC2 instances builds on the familiar Amazon EC2 isolation frameworks such as AWS IAM, security groups, and Amazon VPC. AWS customers maintain control over the Docker daemon itself, the OS, and the underlying EC2 instance when using Amazon ECS. Customers also have control via the AWS deployment and management service family as well as through native configuration via the tooling of their choice. This operational model enables customers to leverage the isolation capabilities of Docker containers and other software isolation frameworks (such as iptables, SELinux, and AppArmor), in conjunction with the underlying AWS security controls, to meet their particular

security, risk, and compliance requirements. For example, customers can provision EC2 instances with different configurations to satisfy security, segregation, or functional requirements. Amazon ECS helps scale this approach through the concept of clusters, which act as a placement boundary for the execution of any given task definition. This relationship enables customers to provision different sets of security configurations and isolation frameworks (such as assignment to different VPCs or EC2 instances) seamlessly and assign them to different sets of task definitions, even as the number of running tasks and container instances scales elastically.

4.11 Software Updates

1) Automated tools will scan for available patches and patch levels, which will be reviewed as specified in the patch application guidelines.

2) Manual scans and reviews will be conducted on systems for which automated tools are not available.

3) An informal risk assessment will be performed within two business days of the receipt of notification of patches. If a determination regarding the applicability of the patch or mitigating controls cannot be made at that time, a formal risk assessment will begin.

4) Vendor-supplied patch documentation will be reviewed to assure compatibility with all system components prior to being applied.

5) Where possible, patches will be successfully tested on nonproduction systems installed with the majority of critical applications/services prior to being loaded on production systems.

6) Successful backups of mission-critical systems will be verified prior to installation of patches and a mechanism for reverting to the patch levels in effect prior to patching will be identified.

7) Patches will be applied during an authorized maintenance window in cases where the patch application will cause a service interruption for mission-critical systems.

8) Patches will be prioritized and applied in accordance with patch application guidelines.

9) Logs will be maintained for all system categories (servers, secure desktops, switches, etc.) indicating which devices have been patched. System logs help record the status of systems and provide continuity among

administrators. The log may be in paper or electronic form. Information to be recorded will include but is not limited to: date of action, administrator's name, patches and patch numbers that were installed, problems encountered, and the system administrator's (SA's) remarks.

10) In the event that a system must be reloaded, all relevant data on the current OS and patch level will be recorded. The system should be brought back to the patch levels in effect before reloading.

11) In the event that a patch will not be applied due to incompatibility or risk assumption, precautions to mitigate the risk of exploitation to the network will be implemented and documented in the log.

4.12 Software IT Administrator Responsibilities

1) SITCORE IT staff are responsible for ensuring that information resources are maintained in compliance with Information Assurance Support Environment (IASE) patch-management policies and procedures.

2) Administrators of systems not managed by IT staff are responsible for ensuring that their systems are maintained in compliance with IASE patch-management policies and procedures (e.g., departmental servers and utility devices).

3) The IT Security Administrator is responsible for auditing information systems to ensure that they comply with IASE patch management policies and procedures.

4.13 Backup Procedures

Various types of cloud-computing backup services are used to protect data in the event of data loss or corruption. The data consist of customer data as well as services and applications including instances within SITCORE.

4.14 EBS Snapshot-Based Protection

When services are running in Amazon EC2, compute instances can use Amazon EBS volumes to store and access primary data. You can use this block storage for structured data such as DBs or unstructured data such as files in a file system on the volume.

Amazon EBS provides the ability to create snapshots (backups) of any Amazon EBS volume. It takes a copy of the volume and places it in Amazon S3, where it is

stored redundantly in multiple availability zones. The first snapshot is a full copy of the volume; ongoing snapshots store incremental block-level changes only.

This is a fast and reliable way to restore full volume data. If you only need a partial restore, you can attach the volume to the running instance under a different device name, mount it, and then use OS copy commands to copy the data from the backup volume to the production volume.

Amazon EBS snapshots can also be copied between AWS regions using the Amazon EBS snapshot copy capability available in the console or from the command line, as described in the Amazon EC2 user's guide.[7] You can use this feature to store your backup in another region without having to manage the underlying replication technology.

4.15 Using Amazon RDS for Backups

Amazon RDS includes features for automating DB backups. Amazon RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance, not just individual DBs.

Amazon RDS provides two different methods for backing up and restoring your DB instances:

1) **Automated backups** enable point-in-time recovery of your DB instance. Automated backups are turned on by default when you create a new DB instance. Amazon RDS performs a full daily backup of your data during a window that you define when you create the DB instance. You can configure a retention period of up to 35 d for the automated backup. Amazon RDS uses these periodic data backups in conjunction with your transaction logs to enable you to restore your DB instance to any second during your retention period, up to the LatestRestorableTime (typically, the last 5 min). To find the latest restorable time for your DB instances, you can use the DescribeDBInstances API call or look on the Description tab for the DB in the Amazon RDS console.

When you initiate a point-in-time recovery, transaction logs are applied to the most appropriate daily backup to restore your DB instance to the time you requested.

2) **DB snapshots** are user-initiated backups that enable you to back up your DB instance to a known state as frequently as you like and then restore to that state at any time. You can use the Amazon RDS console or the CreateDBSnapshot API call to create DB snapshots. These snapshots have

unlimited retention. They are kept until you use the console or the DeleteDBSnapshot API call to explicitly delete them.

When you restore a DB to a point in time or from a DB snapshot, a new DB instance with a new endpoint will be created. In this way, you can create multiple DB instances from a specific DB snapshot or point in time.

4.16 Using AMI to Back Up EC2 Instances

AWS stores system images in what are called Amazon Machine Images (AMIs). These images consist of the template for the root volume required to launch an instance. You can use the AWS Management Console or the AWS EC2 create-image CLI command to back up the root volume as an AMI.[8] Figure 7 shows the process of using an AMI to back up and launch an instance in SITCORE.

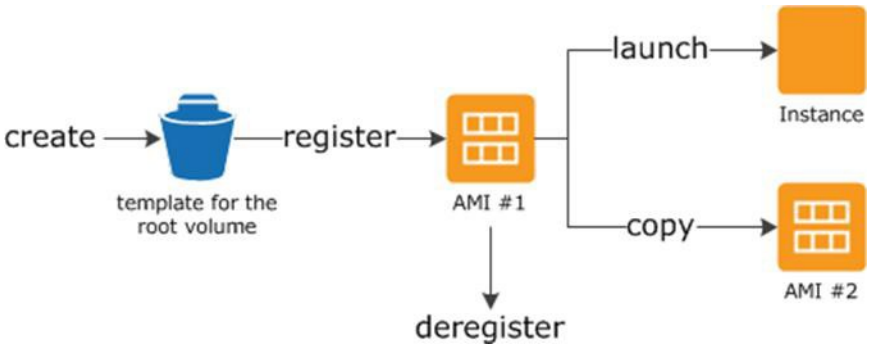

Fig. 7 SITCORE using an AMI to back up and launch an instance

When you register an AMI, it is stored in your account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable.

Figure 8 shows the console using a web instance to create a machine image. After you create an AMI of your Amazon EC2 instance, you can use the AMI to recreate the instance or launch more copies of the instance. You can also copy AMIs from one region to another for application migration or disaster recovery.

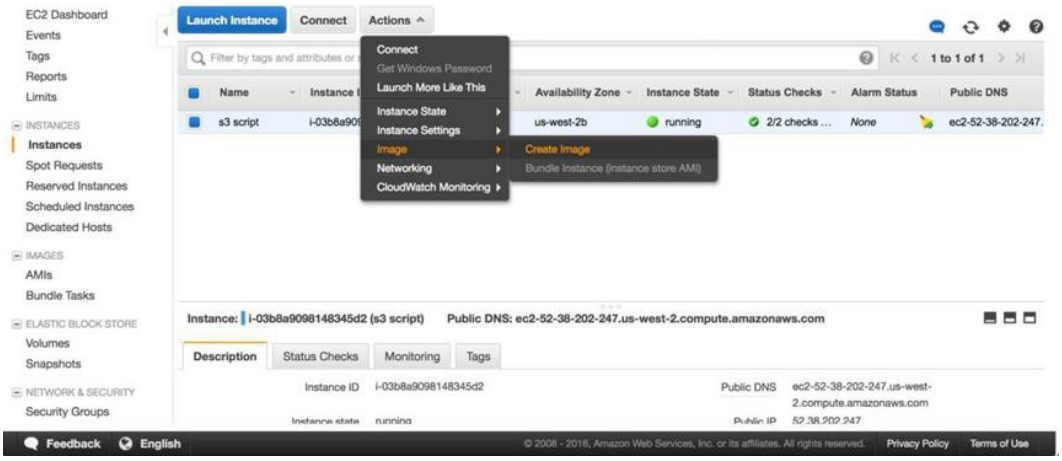

Fig. 8 EC2 console web interface to create a machine image for SITCORE

5. SITCORE SA and Information Assurance Security Officer (IASO)

1) The account request and approval process will be implemented by the SITCORE SA and IASO.
2) The enclave SA/IASO will maintain a record of all hardware and software installed on the instances within SITCORE. The record will be updated on a monthly basis and validated annually.
3) None of the computing devices will be connected to the ARL Enterprise Network.
4) The SITCORE SA will coordinate with the ARL–IAM Scan Team and coordinate scanning of SITCORE devices. This coordination ensures that scans are timely but do not disrupt ongoing experiments.
5) The SITCORE SA will document network topology.

5.1 IAM

The IAM is responsible for the following:

1) Ensures vulnerability compliance via scans.
2) Coordinates with ARL's Computer Network Defense Service Provider to securely configure the DNS border router.
3) Reviews and approves Automated User Permits.
4) Authorizes port access for ARL users for services such as Putty via the port authorization process.

5.2 ARL Networking

There is no ARL networking with SITCORE.

6. Conclusions

ARL is presently developing a cloud-based system to support ARL's SITCORE and Open Campus initiatives. SITCORE provides a virtual collaborative environment testbed for sensor data and information fusion to support collaboration with other government agencies and coalition partners, and provides access to tools, algorithms, facilities, subject-matter experts, and access to data sources for structured and unstructured data. As part of the process of unifying our collaborative research services and tools, ARL is establishing an operational and environmental concept technical document for purposes of defining our platform for use within the ARL research infrastructure to support research and development. This report will serve as a technical document for any future ATO process and the technical basis for system security planning required to operate under an ATO.

References

1. Bennett K, Robertson J. An overview of the US Army Research Laboratory's Sensor Information Testbed Collaborative Research Environment (SITCORE) and Automated Online Data Repository (AODR) capabilities. Proc SPIE. 2017;10190:1019016.
2. Bennett K, Robertson J. Cloud-based security architecture supporting Army Research Laboratory's collaborative research environments. Proc SPIE. 2018;10635:106350G.
3. Bennett K, Robertson J. Advances in the design, development, and deployment of the US Army Research Laboratory (ARL) multimodal signatures database. Proc SPIE. 2011;8040:804009.
4. Bennett K, Robertson J. The US Army Research Laboratory (ARL) Multimodal Signature Database (MMSDB) advanced data storage solutions and security of data over the web. Proc SPIE. 2012;8382:838203.
5. NIST Special Publication 800-53A, Revision 4. Assessing security and privacy controls in federal information systems and organizations. Gaithersburg (MD): National Institute of Standards and Technology; 2015 Jan 29.
6. Bennett K, Robertson J. Security in the cloud: understanding your responsibility. Proc SPIE. 2019;11011:1101106.
7. Amazon EC2. Seattle (WA): Amazon Web Services Inc; c2019 [accessed 2019 Sep 25]. http://aws.amazon.com/ec2/.
8. Copying an Amazon EBS snapshot. Seattle (WA): Amazon Web Services Inc; c2019 [accessed 2019 Sep 25]. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-copy-snapshot.html.

List of Symbols, Abbreviations, and Acronyms

ALB	Application Load Balancer
AMI	Amazon Machine Image
API	application program interfaces
ARL	Army Research Laboratory
ARN	Amazon Resource Name
ATO	Authority to Operate
AWS	Amazon Web Services
CAC	common access card
CLI	Command Line Interface
CPU	central processing unit
CSP	Cloud Service Provider
DB	database
DNS	Domain Name System
DOD	Department of Defense
EAMS-A	Enterprise Access Management Service–Army
EBS	Elastic Block Storage
EC2	Elastic Compute Cloud
ECA	External Certification Authority
ECS	Elastic Container Service
ELB	Elastic Load Balancer
FedRAMP	Federal Risk and Authorization Management Program
HTTP	HyperText Transfer Protocol
HTTPS	HyperText Transfer Protocol Secure
IAM	Identity and Access Management
IASE	Information Assurance Support Environment
IASO	Information Assurance Security Officer

IP	Internet Protocol
IT	information technology
MMSDB	Multimodal Signatures Database
NA	not applicable
NATO	North Atlantic Treaty Organization
OS	operating system
RDS	Relational Database Service
RSA	Rivest–Shamir–Adleman
S3	Simple Storage Service
SA	system administrator
SITCORE	Sensor Information Testbed Collaborative Research Environment
SSH	secure shell
SSL	secure socket layer
SSO	single sign on
URL	Uniform Resource Locator
VPC	Virtual Private Cloud

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	CCDC ARL FCDD RLD CL TECH LIB
2 (PDF)	CCDC ARL FCDD RLS SA K BENNETT J RYDER