



Competition Rules: Preliminary Event 1

Version 1

November 14, 2023



Defense Advanced Research Projects Agency

Biological Technologies Office

675 North Randolph Street

Arlington, VA 22203-2114

TriageChallenge@darpa.mil

Distribution Statement 'A' (Approved for Public Release, Distribution Unlimited)

1 Contents

2	INTRODUCTION	3
3	OVERVIEW	3
4	DTC SCHEDULE OVERVIEW	4
5	PRIZES AND FUNDING	6
6	QUALIFICATIONS	7
6.1	TEAM QUALIFICATION	7
6.2	EVENT QUALIFICATION.....	8
7	DTC TECHNICAL WORKSHOPS	8
8	HUMAN SUBJECTS RESEARCH (HSR)	9
8.1	HANDLING OF DARPA-PROVIDED DATA	9
9	PRIMARY TRIAGE: SYSTEMS COMPETITION RULES	10
9.1	SYSTEMS - ILLUSTRATIVE SCENARIO	10
9.2	SYSTEMS - TECHNICAL CHALLENGE ELEMENTS	10
9.3	SYSTEMS - COMPETITION COURSES	11
9.1	SYSTEMS - EVENT OPERATIONS	12
9.2	SYSTEMS - PERSONNEL GUIDELINES	13
9.3	SYSTEMS - PRELIMINARY EVENT COURSE	16
9.4	SYSTEMS - GUIDELINES	17
9.5	SYSTEMS - SCORING CRITERIA	21
10	PRIMARY TRIAGE: VIRTUAL COMPETITION RULES	23
10.1	VIRTUAL TESTBED.....	23
10.2	VIRTUAL - TECHNICAL CHALLENGE ELEMENTS	25
10.3	VIRTUAL - SCORED EVENT SUBMISSIONS	26
10.4	VIRTUAL - PRELIMINARY EVENT SCENARIOS	28
10.5	VIRTUAL - COMMUNICATIONS AND SCORE REPORTING	29
10.6	VIRTUAL - SCORING CRITERIA.....	30
11	SECONDARY TRIAGE: DATA COMPETITION RULES	32
11.1	DATA - ILLUSTRATIVE SCENARIO	32
11.2	DATA - TECHNICAL CHALLENGE ELEMENTS	32
11.3	DATA TYPES	33
11.4	DATA - SCORED EVENT SUBMISSIONS	33
11.5	DATA - PRELIMINARY EVENT SCENARIOS	35
11.6	DATA - SCORING CRITERIA	37
12	APPENDIX 1 DTC GLOSSARY	40

2 Introduction

This document describes the Preliminary Competition Rules of the DARPA Triage Challenge (DTC). This document is the initial release and revisions prior to the first competition event are expected. The intent of this document is to provide participants guidance on competition design and scoring objectives to inform their development efforts in preparation for the first competition event. This document is subject to change and may be superseded by later versions. The latest official versions of all documents are posted on the [DARPA Triage Challenge Website](http://triagechallenge.darpa.mil) (triagechallenge.darpa.mil) and the [DARPA Triage Challenge Community Forum](#).

DARPA intends to release a draft of the Competition Rules no later than nine months before each Challenge Event. The final version of the Competition Rules will be released no later than three months prior to each respective event.

The DARPA Triage Challenge Chief Judge has the final authority to make any decisions related to the rules or scoring. All decisions made by the Chief Judge are final.

The main goal of the DARPA Triage Challenge is to inspire development of scalable, timely, and accurate capture of novel injury signatures to enhance triage decision-making in austere, complex, and mass-casualty settings. The challenge elements and the competition structure itself are intended to address the additional goal of increasing the diversity, versatility, cost-effectiveness, and robustness of relevant technologies and systems capable of addressing the myriad needs of a wide range of mass casualty incidents (MCIs) rather than single-purpose or specifically tailored solutions. The third goal of the DARPA Triage Challenge is to establish a collaborative community by bringing together multi-disciplinary teams and cross-cutting approaches across disparate fields to address the autonomy, perception, and diagnostic needs of the medical triage community.

3 Overview

Under the authority of 10 U.S.C. §4025 to stimulate innovations using prize competition, the DARPA Triage Challenge will use a series of competition events to drive breakthrough innovations in the identification of physiological features (“signatures”¹) of injury. These new signatures will help medical responders perform scalable, timely, and accurate triage. Of particular interest are MCIs, in both civilian and military settings, when medical resources are limited relative to the need.

The DARPA Triage Challenge’s long-term vision is 1) an initial, or primary stage of MCI triage supported by sensors on stand-off platforms, such as uncrewed aircraft vehicles (UAVs) or uncrewed ground vehicles (UGVs), and algorithms that analyze sensor data in real-time to identify casualties for urgent hands-on evaluation by medical personnel; followed by 2) a secondary stage, after the most urgent casualties have been treated, supported by non-invasive sensors placed on casualties and algorithms that analyze sensor data in real-time to predict the need for life-saving interventions (LSIs) by medical personnel. Injury information provided by these sensors in primary and secondary triage could be integrated with other information about the scene to accumulate evidence about the injury mechanism and characteristics in order to enhance overall situational awareness, and to focus further physiological interventions.

To advance progress towards this vision, the DARPA Triage Challenge aims to bring together multi-disciplinary teams and industries that will identify physiological signatures and develop sensor and

¹ Patterns in sensor data that reflect or predict injuries of high importance for triage assessments.

algorithm strategies for complex MCI settings. Teams participating in the DARPA Triage Challenge will be tasked with developing and demonstrating strategies for capturing high-value signatures for either primary or secondary triage, or for both. While aspects of the DARPA Triage Challenge involve sensors and sensor-delivery platforms, the priority is the development of physiological signatures and models to detect them, not the development of new sensor or platform technology.

4 DARPA Triage Challenge Schedule Overview

The DARPA Triage Challenge is a 3-year effort with 3 sequential 12-month phases for Primary Triage (Systems and Virtual Competitions) and Secondary Triage (Data Competition) in parallel, each culminating in a challenge event (Figure 1; see the [DTC website](#) for competition details). In each phase, competitors will develop signatures and detection and analysis strategies for Primary and/or Secondary Triage. DARPA will host two competition events in each phase; a workshop and a challenge event. Competition events will become progressively more difficult and realistic from Phase 1 to Phase 3.

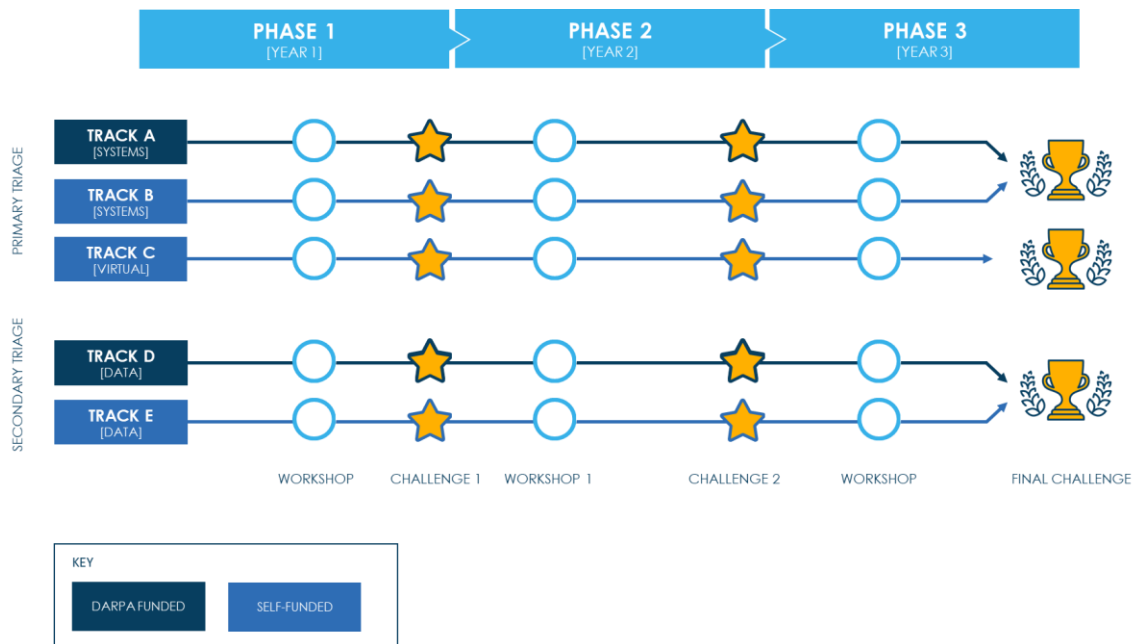


Figure 1 - Program structure and schedule for the DTC.

The workshops will provide an opportunity for practice runs for all tracks and an opportunity for Systems competition teams to collect data from physical simulations of scenarios similar to the end-of-phase challenge event.

Table 1 provides additional information on schedule and format of Competition events and workshops.

Systems Competition - Tracks A and B			
Event	Format	Est. Duration	Date
Year 1			
Challenge Kick-off	In person	2 days	Nov 6-7, 2023
Workshop - Month 8	In person	6 days	Summer 2024
Challenge 1 - Month 12 <i>Evaluations / runs</i>	In person	7 days	Fall 2024
Challenge 1 - Month 12 <i>Awards /lessons-learned session</i>	Hybrid	1 day	Fall 2024
Year 2			
Workshop - Month 4	In person	6 days	Spring 2025
Challenge 2 - Month 12 <i>Evaluations / runs</i>	In person	7 days	Fall 2025
Challenge 2 - Month 12 <i>Awards /lessons-learned session</i>	Hybrid	1 day	Fall 2025
Year 3			
Workshop - Month 4	In person	5 days	Spring 2026
Final Challenge - Month 11 <i>Preliminary Rounds</i>	In person	7 days	Fall 2026
Final Challenge - Month 11 <i>Finalists only - Runs and Awards</i>	In person	1 day	Fall 2026

Virtual Competition and Data Competition - Tracks C, D and E			
Event	Location	Est. Duration	Date
Year 1			
Challenge Kick-off	Hybrid	2 days	Nov 6-7, 2023
Workshop - Month 8	Hybrid	1 day	Summer 2024
Challenge 1 - Month 12 <i>Evaluations / runs</i>	Virtual	TBD	Fall 2024
Challenge 1 - Month 12 <i>Awards /lessons-learned session</i>	Hybrid	1 day	Fall 2024
Year 2			
Workshop - Month 4	Hybrid	1 day	Spring 2025
Challenge 2 - Month 12 <i>Evaluations / runs</i>	Virtual	TBD	Fall 2025
Challenge 2 - Month 12 <i>Awards /lessons-learned session</i>	Hybrid	1 day	Fall 2025
Year 3			
Workshop - Month 4	Hybrid	1 day	Spring 2026

Final Challenge - Month 11 <i>Preliminary Rounds</i>	Virtual	TBD	Fall 2026
Final Challenge - Month 11 <i>Finalists only - Runs and Awards</i>	In person	1 day	Fall 2026

Table 1- Schedule of DARPA-organized Challenge events and workshops. *Note: DARPA-funded teams must attend all workshops in person. It is highly recommended that self-funded Systems teams also attend the workshops in person. For the Challenge events all Systems teams must attend in person.

5 Prizes and Funding

Teams are encouraged to pursue high-risk, high-reward approaches to meet and exceed the objectives of the Challenge Events. Monetary prizes will be awarded for the Systems, Virtual, and Data Competitions at each of the Challenge Events (Table 2).

Domain	Tracks	Phase 1:	Phase 2:	Phase 3 Finals:
Systems Competition	Track A DARPA-Funded	<i>(not eligible)</i>	<i>(not eligible)</i>	Up to \$2.5M prize pool
	Track B Self-Funded	Up to \$200k prize pool	Up to \$500k prize pool	
Virtual Competition	Track C Self-Funded	Up to \$100k prize pool	Up to \$500k prize pool	Up to \$1M prize pool
Data Competition	Track D DARPA-Funded	<i>(not eligible)</i>	<i>(not eligible)</i>	Up to \$1.5M prize pool
	Track E Self-Funded	Up to \$200k prize pool	Up to \$500k prize pool	

Table 2 - Prize structure for the three Challenge Events

DARPA-Funded Teams

DARPA-funded teams (Tracks A and D) are only eligible for the prizes in the Final Events (selection for DARPA-funded team has closed). The Government's obligation for prizes under DTC is subject to the availability of appropriated funds from which payment for prize purposes can be made. No legal liability on the part of the Government for any payment of prizes may arise unless appropriated funds are available to DARPA for such purposes.

Self-Funded Teams

Self-funded teams (Tracks B, C, and E) are eligible for prizes in all of the Challenge Events.
Track B Prizes and Funding: The Phase 1 and Phase 2 prizes for Track B will be awarded to the best performing self-funded Systems Teams in each event, provided that the teams finish in the top 5 overall (including DARPA-funded Track A) teams. High-performing Track B teams are also eligible to become a DARPA-funded Track A team in Phase 2 and/or 3. The Government's obligation for prizes under DARPA Triage Challenge is subject to the availability of appropriated funds from which payment for prize purposes can be made. No legal liability on the part of the Government for any payment of prizes may arise unless appropriated funds are available to DARPA for such purposes.

Track C Prizes and Funding: The Phase 1 and Phase 2 prizes for Track C will be awarded to the best performing self-funded Virtual Teams. The Government's obligation for prizes under DARPA Triage Challenge is subject to the availability of appropriated funds from which payment for prize purposes can be made. No legal liability on the part of the Government for any payment of prizes may arise unless appropriated funds are available to DARPA for such purposes.

Track E Prizes and Funding: The Phase 1 and Phase 2 prizes for Track E will be awarded to the best performing self-funded Data Teams, provided that the teams finish in the top 5 overall (including DARPA-funded Track D) teams. High-performing Track E teams are also eligible to become a DARPA-funded Track D team in Phase 2 and/or 3. The Government's obligation for prizes under DARPA Triage Challenge is subject to the availability of appropriated funds from which payment for prize purposes can be made. No legal liability on the part of the Government for any payment of prizes may arise unless appropriated funds are available to DARPA for such purposes.

To be eligible for prizes, teams must first be registered in the team qualification portal. The award process requires recipients to furnish information that may trace or identify recipients either individually or as an organization (e.g., Social Security Number or Tax Identification Number). The primary contact of each registered team is responsible for providing the award information necessary for prize disbursement. DARPA will reach out by email to the primary contact of each registered team to either confirm their vendor status or request the required forms (e.g., SF-3881 or PIF). DARPA is not responsible for disbursement of prizes to any team members other than the primary contact/organization.

At the end of each competition event, teams will be invited to discuss their technical approaches and lessons learned in a townhall-style hotwash. The extent of technical details shared does not need to exceed data agreements established upon qualification.

6 Qualifications

Prospective DTC competitors must demonstrate track-appropriate performance capabilities to be eligible to participate in DARPA Triage Challenge. All teams in all three competitions (Primary Triage Systems tracks, Primary Triage Virtual track, and Secondary Triage Data tracks; see the [DTC website](#) for track details) must complete two types of qualification: a Team Qualification at the beginning of each phase, and a later event-specific Event Qualifications for each Workshop and Challenge Event. Successful Team Qualification is a prerequisite to Event Qualifications in the same phase.

The event qualification deadline is 100 calendar days before each event. The initial *DTC Event Qualification Guide* is expected to be released by February 18th, 2024. The *DTC Event Qualification Guide* will continue to be updated for each event. The latest revision will be posted on the [DTC Challenge Website](#) and [DTC Community Forum](#).

6.1 Team Qualification

Teams must qualify for DARPA Triage Challenge competition events during the designated qualification window by completing the *Team Qualification* form on the [DTC Team Portal](#). Team Qualification submissions will be accepted on a rolling basis but must be submitted by the deadline (3). Team qualification is required to receive access to datasets and prior to event-specific enrollment.

Team Qualification Windows by Phase

Distribution Statement 'A' (Approved for Public Release, Distribution Unlimited)

Phase 1	9/1/2023 - 11/13/2023
Phase 2	Fall 2024
Phase 3	Fall 2025

Table 3 – Team qualification schedule.

6.2 Event Qualification

Prospective teams are required to demonstrate baseline performance and utility capabilities (e.g., safety measures for the Systems Competition, simulator usage for the Virtual Competition, and algorithm capability for the Data Competition), to be eligible to participate in events. **All** teams (DARPA-funded and self-funded) in all competitions (Systems, Virtual and Data) must qualify for each event including the DTC workshops, Preliminary Events (i.e. Phase 1 and Phase 2 Challenge Events), and Final Event.

The latest revision of the *DTC Event Qualification Guide* will be posted on the DARPA Triage Challenge Website and DTC Discourse Community Forum. Event Qualification submissions will be accepted on a rolling basis but must be submitted by the deadline, approximately 100 calendar days before each event, to be eligible to participate in the event (Table 4). The specific qualification deadlines for each event are provided in the *DTC Event Qualification Guide*.

Failing a previous qualification attempt does not preclude a team from resubmitting a revised qualification submission within the qualification deadlines for any given event. DARPA may adjust the qualification rules for each event and may choose to award qualification waivers for teams that have successfully participated in a prior Workshop or Challenge Event.

DARPA reserves the right to disqualify any team that is found to violate either the rules or applicable laws and regulations.

Event	Event Qualification	Event Date
Workshop 1	3/5/2024	Summer 2024
Challenge 1	Summer 2024	Fall 2024
Workshop 2	Winter 2024-2025	Winter 2024-2025
Challenge 2	Summer 2025	Fall 2025
Workshop 3	Winter 2025-2026	Winter 2025-2026
Challenge 3	Summer 2026	Fall 2026

Table 4 – Event qualification schedule.

7 DARPA Triage Challenge Technical Workshops

DARPA encourages vibrant information exchange and collaborative interactions among all DARPA Triage Challenge participants, to include DARPA technical staff, independent verification and validation (IV&V) teams, representatives from competitor teams, infrastructure developers, and other government partners. To that end, DARPA will host a workshop in each phase which will offer a forum for community building and cross-pollination of technical ideas and approaches as well as an opportunity for testing in the Systems Competition.

In each phase (8 months into Phase 1, 4 months into Phases 2 and 3) DARPA will host a multi-day hybrid workshop. This will include live practice sessions for Systems Competition competitors to test their systems on simulated casualty scenes similar to the next challenge event and virtual practice sessions for Virtual Competition teams to test their algorithms in simulated casualty scenes similar to the next

challenge event. The practice sessions will be followed by a ‘lessons learned’ discussion for all tracks and an opportunity to discuss real-world needs with Government partners.

At the workshops, teams will have opportunities to rehearse their runs, confirm integration with the DARPA instrumentation and scoring systems, and inform their development efforts. Runs at the workshops are not officially scored, but teams are encouraged to operate according to the Competition Rules to prepare for the Challenge events. In-person attendance at workshop events is required for all DARPA-funded teams. Self-funded teams may choose to attend virtually or in person, although Self-funded teams on the Systems track are **strongly** encouraged to attend in person.

8 Human Subjects Research (HSR)

For the Primary Triage Competition, Systems teams must be included in the IV&V Team’s Institutional Review Board (IRB) protocol through a DoD Institutional Agreement for Institutional Review Board Review (IAIR) to access training data collected by the IV&V team and to collect data at DTC workshops and challenge events. For the Secondary Triage Competition, use of training data provided by DARPA does not constitute HSR, and competitors do not need to obtain IRB approval to use these data. For both Primary and Secondary Triage Competitions, DARPA-funded competitors require DARPA approval for the collection or use of any other human subject data. **Self-funded teams are prohibited from the collection or use of any other human subject data as part of their involvement in the DARPA Triage Challenge, beyond data and data-collection opportunities provided by DARPA, because DARPA HSR supervision is not feasible for teams not under DARPA contract.** Self-funded teams should carefully consider this limitation and should take this into account in their technical approach, leveraging other strategies as appropriate (*e.g.*, simulations).

DoD Definition of Human Subjects Research (HSR)

The term “human subject” can be applied to research efforts that meet EITHER of the following criteria:

A living individual about whom an investigator (whether professional or student) conducting research:

- Obtains information or biospecimens through intervention or interaction with the individual, and uses, studies, or analyzes the information or biospecimens; or
- Obtains, uses, studies, analyzes, or generates identifiable private information, personally identifiable information, or identifiable biospecimens.

Human Subjects Research involves:

- Activities that include both a systematic investigation designed to develop or contribute to generalizable knowledge and involve a living individual about whom an investigator conducting research obtains information or biospecimens through intervention or interaction with the individual, or identifiable private information, or biospecimens.

8.1 Handling of DARPA-provided data

Primary Triage:

Primary triage datasets are owned by the Army and developed by its Telemedicine & Advanced Technology Research Center (TATRC), and shared with DARPA under appropriate authorities, exclusively for research purposes (including DTC). The TATRC datasets entrusted to DARPA have been intentionally de-identified to ensure—to the greatest extent practicable—that there is no reasonable basis to believe that the data could be used to trace a specific identity or present a risk of harm to any individual. However, TATRC datasets may incidentally or unintentionally contain sensitive information and images (including facial imagery). Therefore, as previously acknowledged in the DTC Qualification

process, competitors agree they will *not* attempt to re-identify, share, or re-use Army/TATRC data as provided by DARPA.

Secondary Triage:

Secondary triage datasets are provided by the DARPA Research Infrastructure for Trauma with Medical Observations (RITMO) research program for use during DTC. DARPA’s mission requirement and intent are to safeguard privacy and civil liberties and the RITMO datasets have been intentionally de-identified to ensure—to the greatest extent practicable—that there is no reasonable basis to believe that the data could be used to trace a specific identity or present a risk of harm to any individual. Therefore, as previously acknowledged in the DTC Qualification process, competitors agree that they will *not* attempt to re-identify data in the RITMO datasets, nor attempt to download or share the RITMO datasets.

9 Primary Triage: Systems Competition Rules

9.1 Systems - Illustrative Scenario

The notional DARPA Triage Challenge primary triage setting is the first few minutes of an MCI where the number of casualties and/or the environment likely would preclude a timely initial assessment of each casualty by first responders.

The objective of the Primary triage competitions is to detect and identify physiological signatures of injury derived from data captured by stand-off sensors to enable early prioritization of casualties, allowing medical care professionals to quickly focus on the most urgent casualties. Competitors will develop algorithms that detect those signatures in real-time from stand-off sensors on robotic mobility platforms (e.g., UAVs, UGVs) to provide decision support appropriate for austere and complex pre-hospital settings. Of particular interest are signatures of acutely life-threatening conditions that medics are trained and equipped to treat during primary triage, such as hemorrhage and airway injuries.

Challenge events for System competitors will be physical simulations of casualty scenarios. Although the setting and complexity of challenge events will vary over the course of DTC, the following features are expected to be maintained across events. Each competitor will have access to the same scenario and no two teams will operate on the same location simultaneously. Competitors will have only general information on the setting beforehand—for example, that it is a battlefield scenario, or a collapsed building following an earthquake. There will be actors and manikins exhibiting simulated injuries of varied type and severity (subject to the limitations of what can be simulated).

Competitor systems with stand-off sensors, robotic mobility platforms, and algorithms will need to autonomously process sensor data and provide real-time casualty identification and injury assessment. No part of a competitor’s system may touch a casualty or manipulate the scene (e.g., clear rubble). Each scenario will have a time limit, with no scenario expected to have a duration greater than 30 minutes in Challenge 1.

9.2 Systems - Technical Challenge Elements

The Challenge competition courses will be designed to assess performance across various challenge elements, including: Degraded sensing, Obscuring obstacles, Terrain Obstacles, Dynamic Obstacles, and Dynamic Casualties. The challenge elements are expected to become progressively more difficult from Phase 1 to Phase 3.

1. *Degraded Sensing*: The courses are expected to include elements that range from constrained passages to large open fields, lighted areas to complete darkness, and wet to dusty conditions. Sensors will need to have the dynamic range to reliably operate in these environments. Dust, fog, mist, smoke, talking, flashing light, heat spots, and loud background audio effects are within scope of this challenge element. Extreme temperatures, fire, and hazardous materials are not expected to be within scope.
2. *Obscuring obstacles*: Casualties may be fully visible to partially obscured to completely obscured, such as buried under a shallow layer of rubble. Sensor modalities capable of penetrating rubble will have an advantage in such situations. Stand-off sensor access to skin may be possible but cannot be assumed. Casualties may also be grouped with limbs overlapping, or may be interacting with live responders.
3. *Terrain Obstacles*: Systems will be required to demonstrate robustness in navigating a range of terrain features and obstacles. Terrain elements and obstacles may include constrained passages, large drops/climbs, inclines, and rubble. The environments may include natural or human-made materials; structured or unstructured clutter; and intact or collapsed structures and debris.
4. *Dynamic Obstacles*: Live responders, “walking wounded”, or other physical changes to the environment will test the agility of the system autonomy to identify and assess casualties.
5. *Dynamic Casualties*: Some treatable injuries may rapidly be fatal, so taking too long to find and assess casualties may result in the window for effective LSI being missed. While competitors are not expected to re-evaluate casualties for changes in status, casualties who are not evaluated within an appropriate timescale may have a change in status (for example, progression of untreated hemorrhage or airway injury).
6. *Endurance Limits*: It is expected that individual scenarios will run between 15-30 minutes. Teams will be permitted to replace batteries during their run, but teams should consider the implications of returning to the original launch location and redeploying their systems.

9.3 Systems - Competition Courses

There are two themes for Challenge 1: a plane crash and a post-battle environment. The three course scenarios in year 1 will focus on these themes. All courses are expected to be located outdoors during daylight where the ground will be pavement, grass or gravel. Obstacles such as small piles of rubble, buildings and vehicles may be placed on the scene. Casualties (manikins and actors) will be located on the ground in various locations. Teams will take turns to deploy their systems on scene and the scenes are reset between runs. In the event of high wind or rain, runs may be delayed. Courses are expected to be 100 ft x 100 ft x 100 ft or smaller.

Figure 2 shows a notional workflow and communications plan for the competition events. The competing team will set up and begin their run in the designated Staging Area. At the beginning of a run, teams will deploy their systems onto the course where they will explore, locate, and triage casualties for no longer than the time limit set for the scenario. Observation data will be transmitted to the team’s Base Station which will, in turn, provide triage reports to the DARPA Command Post (CP) where the reports will be automatically evaluated and scored. The DARPA Command Post will provide score updates back to the team’s Base Station.

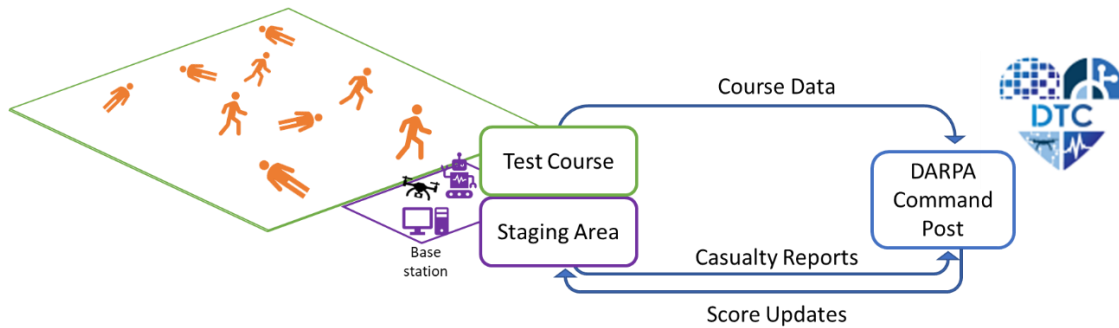


Figure 2 - Course workflow and data transmission

An Interface Control Document (ICD) and reference implementation will detail the mechanism for providing triage reports to DARPA. The latest revision of the *DTC* ICD will always be posted on the DTC Website and DTC Community Forum.

9.1 Systems - Event Operations

9.1.1 Competition Format

Prospective teams are required to demonstrate baseline performance and utility capabilities, as described in Section 0, to qualify for Challenge Event 1. It is anticipated that up to 16 teams may successfully qualify for the event. The event is expected to include three competition days. Qualified teams will be eligible to participate in the event, which will consist of three scored runs 15-30 minutes in duration. The total score for the event will be sum of a team's best two runs.

9.1.2 Staging Area

All systems will be required to start in the Staging Area behind the Starting Gate at the course entrance. No systems will be permitted to operate outside of the competition course boundaries except within the Staging Area. Additional details about the Staging Area including layout and dimensions will be provided at a later time.

9.1.3 Course Access

Systems are allowed to enter, exit, or reenter the competition course at any time within the duration of the run. All human operators and personnel must stay within the Staging Area. No manual physical intervention or entry by any (human) team member on the course will be permitted. A system may only be handled or retrieved if it has crossed back into the Staging Area past the front face of the Starting Gate. Once a system has partially or completely crossed into the Staging Area, team personnel are permitted to handle the systems as long as the personnel stay within the Staging Area and do not pass the front face of the Starting Gate. Only authorized DARPA personnel are allowed to enter the course preceding, during, and following the run.

9.1.4 Run Termination

A scored run terminates upon any of the following conditions:

- Time Expiration: The scored run time expires before another termination criterion is met.
- Run Completion: The deployed systems successfully report on all casualties and exit the course.

Distribution Statement 'A' (Approved for Public Release, Distribution Unlimited)

- **Run Cancellation:** Competition Staff cancels the run due to an external factor such as weather, including lightning, rain, snow, or wind.
- **Emergency Stop:** Competition Staff initiates an emergency stop because of an unsafe condition.
- **By Request:** The Team Lead requests an end to the run.

9.1.5 Terminated Runs

A team may be eligible for an additional attempt if a run is canceled or stopped due to an emergency or external factor outside of the team’s control. The Chief Judge will review eligible cases and determine the course of action. The Chief Judge has the final authority to make any scoring-related decisions.

9.1.6 Score Disputes

Dispute Cards are intended to provide teams a mechanism to submit a formal dispute or request for review by the Chief Judge. The Dispute Card must be completed and delivered by the Team Lead to the relevant Course Official, Team Garage Coordinator, or Chief Judge. The Dispute Card must be submitted within 30 minutes of the completion of the run in question. All submissions will be reviewed by the Chief Judge in a timely manner. All decisions made by the Chief Judge are final.

9.2 Systems - Personnel Guidelines

Teams are permitted up to six (6) personnel in the Staging Area; these personnel are designated Pit Crew. Figure 3 provides a detailed workflow for how data may be shared between the systems, team Base Station, team personnel, and DARPA Command Post. Two categories of data are delineated: Status Data and Derived Data. Status Data includes real-time sensor streams from the deployed systems for the purposes of calibration, system status monitoring, teleoperation, and safety monitoring. Derived Data includes data that has been processed or fused to create derived information from the raw sensor streams. Derived Data specifically includes any casualty reports.

Pit Crew personnel are permitted to assist with operations tasks such as physically deploying the systems, performing repairs, modifying software or firmware, and changing batteries. Pit Crew personnel are only permitted to access limited Status Data. They are not permitted to wirelessly communicate with the deployed systems and are not permitted to access Derived Data or Casualty Reports. Pit Crew may take on one of three additional specialized roles:

- **Safety Officers** are responsible for preserving the safety of personnel and property and are permitted to communicate with the deployed systems solely for safety purposes.
- **Operators** are permitted to communicate with deployed UxV systems, teleoperate deployed UxV systems, activate safety emergency stops, and access Status Data.
- **Human Supervisors** are permitted to communicate with deployed systems, teleoperate deployed systems, activate safety emergency stops, and access Status Data.

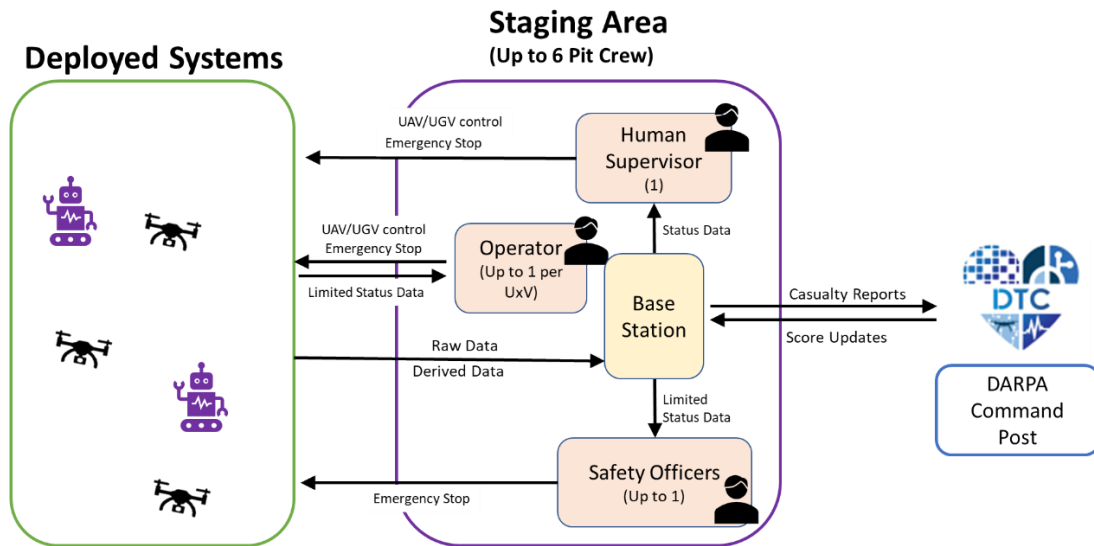


Figure 3 - Data Workflow for the Systems Competition

9.2.1 Pit Crew Personnel

The role of the Pit Crew is to assist with operations tasks such as physically deploying the systems, performing repairs, modifying software or firmware, and changing batteries. Once a team's run has begun, the Pit Crew personnel may not be substituted with other personnel.

The Pit Crew personnel, including Operators, Safety Officers and Human Supervisors, are permitted to verbally communicate without restrictions. The Base Station can also provide limited status data to the Pit Crew via a wired display to support operations tasks such as calibration and completing startup checklists. The Pit Crew is not permitted to directly interface with the Base Station in any way (e.g., toggling between windows via peripherals). The Pit Crew is only permitted to view limited system status data such as battery health, network status, and real-time telemetry.

Pit Crew personnel are permitted to view and access Status Data but are not permitted to view or access Derived Data. Pit Crew personnel are specifically prohibited from viewing or accessing Casualty Reports.

9.2.2 Safety Officers

The role of the Safety Officer is to preserve the safety of personnel and property. Safety Officers are permitted to activate Tier 1 wireless emergency stop transmitters and/or operate remote controls for safety purposes only. Safety Officers may only use wireless communications for emergency stop transmitters and limited system initialization (e.g., arming, initial takeoff).

The team is permitted to have up to one Safety Officer in addition to any Operators.

For aerial systems, the Safety Officer may aid in initial takeoff and hover as long as the system is within the Staging Area and does not intrude into the competition course. However, any further maneuvering of the aerial system must be initiated or controlled by a Human Supervisor or Operator. If the safety of personnel is at risk, the Safety Officer is permitted to take control of the aerial system for the sole purpose of safely landing the system. The Safety Officer is not permitted to take control of a system that has crossed into the competition course except to trigger a Tier 1 emergency stop for the sole purpose of preserving the safety of personnel.

The Safety Officer’s primary role is to preserve the safety of personnel in the Staging Area rather than preserving the safety of the system. If the Safety Officer triggers a Tier 1 emergency stop for a system inside the competition course, the system that is triggered must stay inactive for the remainder of the competition run. A system may only have its Tier 1 emergency stop reset if it was triggered by a Human Supervisor.

Safety Officers are also permitted to perform all the roles of the Pit Crew personnel.

9.2.3 Operators

The role of the Operator is to operate one UxV system during the team’s run. This can include take off, landing, and full movement throughout the course. Teams are allowed one Operator per UxV system for a total of up to 5 operators. Any individual serving in the Operator role must have a valid part 107 license if piloting a UAV.

At this time, DARPA plans to allow Operators only for Challenge Event 1.

9.2.4 Human Supervisor

As the operational scenario suggests, DARPA is interested in approaches that are mostly autonomous without the need for substantive human interventions, and capable of remotely locating and providing assessments of causalities. The team is permitted to have one Human Supervisor per team, Human Supervisors are required to have valid part 107 license if the team has aerial vehicles.

The Human Supervisors are permitted to monitor and manage the communications with their deployed systems. The Human Supervisors are permitted to view, access, and/or analyze Status Data but are not permitted to view or access Derived Data or Casualty Reports. Once a team’s run has begun, the Human Supervisors may not be substituted with other personnel.

Human Supervisors are also permitted to perform all the roles of the Safety Officers, Operators, and Pit Crew personnel.

Role	Responsibilities and Access					
	View Limited Status Data	View Full Status Data	View Derived Data	Trigger Tier 1 EStop	Service UxV	Teleoperation Of UxV
General Pit Crew (Up to 6)	✓				✓	
Pit Crew: Specialized Roles (each general pit crew teammate permitted up to one specialized role)						
Human Supervisor (1)	✓	✓		✓	✓	✓

Operator (1 per UxV up to 5)	✓			✓	✓	✓
Safety Officer (up to 1)	✓			✓	✓	

Table 5 - Roles and Responsibilities of Pit Crew.

9.2.5 Teleoperation

In Phase 1, competitors will be permitted to teleoperate their deployed systems with up to one Operator per UxV system and one Human Supervisor per team. DARPA expects to place additional limits on teleoperation in later phases. By Phase 3, it is expected that only one Human Supervisor will be permitted to teleoperate all of the deployed systems and will be the only team member permitted to communicate with the deployed systems.

Due to the complicated nature of the course, tethers for power, communications, or physical retrieval are not permitted. No manual physical intervention or entry by any (human) team members on the course is permitted. Only authorized DARPA personnel are allowed to enter the course preceding, during, and following the run.

9.3 Systems - Preliminary Event Course

9.3.1 Course Layout

The competition courses are expected to simulate a variety of mass casualty incidents such as a large traffic accident or a collapsed building. The exact course layout will not be known to competitors in advance, and DARPA intends to alter the competition courses to randomize casualty types and casualty locations. More details related to competition courses will be provided closer to the competition events.

9.3.2 Challenge 1 Event Course Challenge Elements

The scale and complexity of the three competition courses in year 1 is expected to vary. The design of the first Challenge Event courses is intended to assess the ability of teams to address the variety of challenging environments presented in a plane crash or post-battle environment. The following subset of challenge elements are anticipated to be present in Year 1.

1. *Obscuring obstacles:* Casualties may be fully visible to partially obscured. Stand-off sensor access to skin may be possible but cannot be assumed (e.g., clothing, makeup, dust).
2. *Terrain Obstacles:* Systems will be required to demonstrate robustness in navigating a range of terrain features and obstacles. Terrain elements and obstacles may include inclines, vehicles, buildings and rubble. The environments may include natural or human-made materials; structured or unstructured clutter; and debris.
3. *Endurance Limits:* It is expected that individual scenarios will run between 15-30 minutes. Teams may be permitted to replace batteries in their staging area during their run, but teams should consider the implications of returning to the original launch location and redeploying their systems.

9.4 Systems - Guidelines

9.4.1 Prohibition on deployment of humans or animals

Teams may choose to deploy a wide variety of systems to complete the course objectives including but not limited to robotic platforms, sensors, and communication components. No humans or animals will be permitted as any part of the deployed systems that enter the competition course.

9.4.2 NDAA Compliance

UAVs must be National Defense Authorization Act (NDAA)-compliant. All teams will need to submit documentation on the NDAA compliance of all UAVs. Below is the definition of prohibited UAVs reproduced with permission from <https://www.diu.mil/blue-uas-policy>.

The National Defense Authorization Act (NDAA) for Fiscal Year 2023 was passed on 23 December, 2022 and Section 817 remains in effect. Section 817 modified portions of FY23 NDAA Sec 848. The bill can be found in its entirety at [Congress.gov](https://www.congress.gov).

Definitions

These definitions are extracted with permission from the Procedures for the Operation or Procurement of Unmanned Aircraft Systems to Implement Section 848 of the NDAA for Fiscal Year 2020, published 2 September, 2021. Terms that were modified by FY23 NDAA Sec 817 have been updated.

Covered UAS: Any UAS and any related equipment that:

- 1. Are manufactured in a covered foreign country or by an entity domiciled in a covered foreign country;*
- 2. Contain critical components, as defined in this document, manufactured in a covered foreign country or by an entity domiciled in a covered foreign country;*
- 3. Use a ground control system or operating software developed in a covered foreign country or by an entity domiciled in a covered foreign country; or*
- 4. Use network connectivity or data storage located in or administered by an entity domiciled in a covered foreign country*

Covered UAS Company: Any of the following:

- 1. Da-Jiang Innovations (or any subsidiary or affiliate of Da-Jiang Innovations)*
- 2. Any entity that produces or provides unmanned aircraft systems and is included on Consolidated Screening List maintained by the International Trade Administration of the Department of Commerce*
- 3. Any entity that produces or provides unmanned aircraft systems and—*
 - 1. is domiciled in a covered foreign country; or*
 - 2. is subject to unmitigated foreign ownership, control or influence by a covered foreign country, as determined by the Secretary of Defense unmitigated foreign ownership, control or influence in accordance with the National Industrial Security Program (or any successor to such program).*

The term "covered foreign country" means the People's Republic of China, the Russian Federation, the Islamic Republic of Iran and the Democratic People's Republic of Korea.

The term "place of manufacture" has the definition provided in FAR 52.225-18, as the "place where an end product is assembled out of components, or otherwise made or processed from raw materials into the finished product that is to be provided to the Government." If a product is disassembled and reassembled, the place of reassembly is not the place of manufacture.

The following are included in the definition of "critical components":

- 1. Flight controller: The combination of embedded software on computing hardware, that issues commands to actuators based on the difference between the desired and actual position of a UAS.*
- 2. Radio: A device that enables communication by packaging, transmitting, and/or receiving modulated signals into or from electromagnetic waves in the radio frequency (RF) spectrum.*
- 3. Data transmission device: Electronic hardware that actively transfers electronic information from one digital system to another.*
- 4. Camera: A device that converts focused light onto a photosensitive sensor for the purpose of recording or transmitting visual images in the form of photographs, film, or video signals.*
- 5. Gimbal: A mechanism, typically consisting of electromechanical actuators and a mechanical frame, which rotates about one or more axes to stabilize and properly orient cameras or other sensors.*
- 6. Ground control system: An electronic mechanism that enables a human operator to transmit data in order to influence the actions of an aerial vehicle remotely.*
- 7. Operating software: A program that directs a computer's basic functions, such as scheduling tasks, executing applications, and controlling peripherals.*
- 8. Network connectivity: The hardware and software required for communication between computers over the internet or other distributed and separately administered systems, for example, through the use of routers, switches, and gateways.*
- 9. Data storage: The collective methods and technologies that capture and retain digital information on electromagnetic, optical, or silicon-based storage media.*

9.4.3 System Constraints

- UGV maximum weight = 200kg
- UAV/UGV maximum diameter = 1.5m
- UAV maximum weight = 9kg
- DARPA is interested in portable systems. Therefore, all system elements must pack down to be carried by a single vehicle (car, sports utility vehicle or pickup truck),
- Platforms must not produce any visible illumination other than what is legally required for UAV flight.

9.4.4 Stand-off distances

- UGVs minimum stand-off distance = 1 meter
- UAVs will have variable minimum stand-off distances based on size
 - Less than 250 g: 1 m
 - Between 250 g and 5 kg: 3 m
 - Greater than 5 kg = 5 m
- UAV maximum altitude = 30 m

9.4.5 Sensor Constraints

Distribution Statement 'A' (Approved for Public Release, Distribution Unlimited)

DARPA expects that multi-modal approaches will be required to improve signature identification and address multiple challenge elements that could degrade the usefulness of any one sensor (e.g., environmental conditions, casualty pose). Various sensor modalities and combinations will be allowed, including but not limited to LIDAR, acoustic, visual, RF, IR, UV, radar, gravity, compass/magnetic, GPS, and chemical.

- Sensors must be capable of detecting the desired signatures from the relevant stand-off distances (see 9.4.4).
- All sensor elements must be skin- and eye-safe.
- Sensors may not physically interact with casualties.
- All audio communication must be autonomous: no human communication is permitted between teams and casualties.

9.4.6 System Cost and Quantity Constraints

Teams are limited to a maximum of 5 deployed autonomous systems.

DARPA is interested in solutions that are cost-effective and attrition-tolerant. Due to the complexity of the environments, teams should expect and plan for some level of failures and/or attrition. While there are currently no limits on the aggregate cost of deployed systems, DARPA may introduce additional constraints as the competition progresses to appropriately incentivize such solutions.

9.4.7 System Retrieval

All systems must begin the run in the Staging Area. It is encouraged but not required for the deployed systems to return to the Staging Area at the end of the run. Any systems that have not autonomously exited the course at the termination of a run will be retrieved by authorized DARPA Competition Staff. The Competition Staff will make their best effort to collect a team's systems after each run. However, if systems are not able to be recovered in a safe or timely manner by the Competition Staff, teams may have to operate without them on subsequent runs. Teams are encouraged to provide audible and visual recovery aids such as flashing LEDs and audible cues to help Competition Staff locate deployed systems or components.

Due to the need for exclusive handling by Competition Staff, several safety measures will be required for robotic platforms. These include but are not limited to a DARPA-specified transponder and a DARPA-approved emergency stop.

9.4.8 Emergency Stop

The emergency stop (E-Stop) requirements are designed to ensure the safety of personnel, equipment, and the competition course environment. All systems participating in the Systems Competition will utilize a complementary three-tiered emergency stop system.

Tier 1: Team Wireless E-Stop

Teams are required to implement a wireless emergency stop capability as a component of their system's communication architecture. The emergency stop must be able to be triggered from the team's Base Station and/or portable wireless transmitter. The Tier 1 E-Stop transmitter must instruct mobile platforms within effective communication range to come to a halt. The emergency stop procedures implemented on the mobile platforms must, upon receiving a Tier 1 E-Stop trigger, render a platform completely motionless within 30 seconds. The emergency stop must include clear visual feedback of the mobile

platform's safe, halted state (e.g., red LED). The emergency stop capability may be targeted to a specific platform, but should also provide the functionality to rapidly render all platforms safe. A team must be able to render all platforms within communication range completely motionless within 60 seconds.

Tier 2: Recovery Wireless E-Stop

DARPA is considering a tier 2 E-stop. Further details TBD.

Tier 3: On-Platform E-Stop

Teams must integrate at least one emergency stop button on each platform that weighs more than 10 kg. The button must be a red mushroom-capped button at least 25 mm in diameter, with clear markings indicating that it is an emergency stop button. The buttons must latch when triggered and must require a twisting motion to release the latch. The buttons must be completely unobstructed and must be easily accessible by recovery personnel. The emergency stop procedures implemented on the mobile platforms must, upon receiving a Tier 3 E-Stop trigger, render all platforms completely motionless within 5 seconds.

E-Stop Qualification:

In accordance with the DTC Event Qualification Guide document, all teams are required to demonstrate emergency stop compliance to be eligible for participation in the Competition Events. Qualification requires teams to demonstrate fully functional emergency stopping in compliance with all three tiers outlined in this document. Demonstration requirements are outlined in the "Emergency Stop" section of the *DTC Event Qualification Guide*.

Emergency stop functionality and compliance will be verified by DARPA at each official DTC Challenge event. DARPA reserves the right to deny a team's participation in one or more runs if any of their mobile platforms are non-compliant with the emergency stop rule.

9.4.9 Dropped Components

Teams are not permitted to use dropped components and leave-behind peripherals.

9.4.10 Course Alteration

The course may not be willfully altered by any of the deployed systems, including but not limited to digging, burrowing, or intentional degradation or destruction of the environment's walls, floors, ceiling, immobile barriers or obstacles, or other course infrastructure or instrument.

9.4.11 Power Sources

All fuel and power sources will need to be approved by DARPA for use in the competition. Teams may be required to submit safety protocols and DARPA may require additional site-specific approvals which could require significant lead time. Most electric battery sources are expected to be approved. Teams are encouraged to address any potential concerns early.

9.4.12 Competition Networking

Casualty reports will be submitted from the base station over a DARPA-managed wired network. A wireless team's network will be available for teleoperation and data transmission between platforms and the base station. Additional details will be announced closer to the first competition event.

9.4.13 Internet and Cloud Resources

DARPA does not plan to provide or allow the use of internet (www) or cloud connectivity during the runs in the Systems Competition. Access to such resources is often limited in the field and in real-world scenarios following natural disasters. Team personnel in the Staging Area are not permitted to access the internet or make phone calls on any devices (e.g., cell phone, tablet) during the competition run.

9.4.14 Data protection

Teams are required to use encryption for data transmission between their base station and their UxVs. Data transfer from the bases station computers to a computer approved for processing, analysis and/or data storage will be performed using approved transfer methods such as writeable CDs, DVDs, or memory cards, after being scanned for malware. At the end of an event, teams will remove and clear any removable memory in the UxV.

9.5 Systems - Scoring Criteria

The goal of the DARPA Triage Challenge is to develop scalable, timely and accurate capture of novel injury signatures to enhance triage decision-making in austere, complex, and mass-casualty settings.

In the Primary Triage challenge (both Systems and Virtual), teams are evaluated based on accuracy and speed in assessing casualty condition using one or more autonomous platforms and stand-off sensors. Casualties are distributed throughout the competition course in a manner which rewards teams that are able to rapidly find and assess casualties. The nature of the casualties is not known prior to a run by competitors and may be varied from run to run.

Upon identifying a casualty status, the deployed system must report its injury diagnosis and relevant clinical information to the evaluation system for scoring. Reports will be submitted to the DARPA Command Post via the team's Base Station over an Ethernet link. The detailed report format, protocol, and example implementation will be specified in the ICD. The reported casualty conditions are compared against coincident ground truth data for scoring.

Upon submitting a casualty report, teams receive a response including confirmation of receipt and the score update. Score updates during a run are expected to happen immediately after report receipt. Additional details will be provided in the ICD available on the [DTC website](#).

9.5.1 Casualty Localization and Identification

For the phase 1 challenge, each casualty in the scene will have a visual identifier (e.g., QR code) nearby indicating the identity of the casualty. Casualties and their visual identifiers will be positioned such that there is an unambiguous association between the two. The unique casualty identifier will need to be included in the triage report for scoring, as specified in the ICD. In future years, identification and scoring will require the systems to localize casualties relative to a predetermined origin position. Details will be released in future versions of this document.

9.5.2 Casualty Reports

Reports capture information about clinical condition of casualties relevant for triage decision making. The casualty report is comprised of two parts: the first part contains vital signs and time-critical indicators of immediate need, and the second part contains more detailed assessment of injury condition and severity. Teams may report both parts simultaneously, or they may report each part at separate times. Any submitted report, whether partial or full, contains a casualty identifier and timestamp.

In future years, reports will include casualty location and may include additional clinical assessment features. The preliminary list of clinical features in the report relevant for scoring is shown in Table 3. Complete details on report contents and format will be found in the ICD.

9.5.3 Report scoring

A valid casualty report earns the team up to 10 points, with up to 5 additional bonus points for early reporting of time-critical information. Points are determined by comparing report contents to ground truth data from the same time. Categorical fields in the report are awarded points based on whether they match ground truth (GT); numerical fields in the report are awarded points based on whether they are within a predetermined range of the ground truth value.

To incentivize rapid assessment of time-critical information indicating immediate need of medical care, bonus points may be awarded for early casualty reports containing the following fields: Severe Hemorrhage, Respiratory Distress, and vital signs (Heart Rate and Respiratory Rate). Bonus points will be awarded for valid reports received by the DARPA Command Post within an initial “golden window” during a scored run. The duration of the “golden window” will be provided to teams prior to the challenge. Note that the report contents with bonus potential can be reported separately from those that do not have bonus potential, as described in the previous section.

Preliminary scoring criteria and bonus potential for clinical assessment in the casualty report are shown in Table 6. Casualty report definitions and examples will be found in the ICD.

Field	Values	Scoring Criteria
Severe Hemorrhage ^{1,3}	[Present, Absent]	2 if match ground truth (GT) 0 otherwise
Respiratory Distress ^{1,3}	[Present, Absent]	2 if match GT 0 otherwise
Heart Rate ^{1,4}	Integer	1 if within <i>n</i> of GT 0 otherwise
Respiratory Rate ^{1,4}	Integer	1 if within <i>m</i> of GT 0 otherwise
Trauma ²	Head: [Wound, Absent] Torso: [Wound, Absent] Upper Ext.: [Wound, Amputation, Absent] Lower Ext.: [Wound, Amputation, Absent]	2 if all match GT 1 if at least two match GT 0 otherwise
Alertness ²	Ocular: [Open, Closed, Not Testable (NT)] Verbal: [Normal, Abnormal, Absent, NT] Motor: [Normal, Abnormal, Absent, NT]	2 if all match GT 1 if at least two match GT 0 otherwise

Table 6 - Preliminary casualty report clinical assessment with scoring criteria.

¹ Time-critical features comprising the first part of the report.

² Detailed injury and severity assessment comprising the second part of the report.

³ Correct response receives 2x bonus if reported within “golden window”.

⁴ Correct response receives 1.5x bonus if reported within “golden window”.

9.5.4 Time

Time is measured from the start of the run and scoring will be based off of two timestamps: the assessment timestamp and the report timestamp. The assessment timestamp is provided by the deployed system in the report and indicates the time at which the casualty's condition was assessed. This timestamp is used to determine the appropriate ground truth for scoring. The report timestamp is the time when the valid report is received at the DARPA command Post, and it determines any bonus points awarded.

Report limits

To discourage guessing and preserve system bandwidth, the DARPA Command Post will limit the total number of scored reports. Any further reports are rejected. In Phase 1 it is expected that each run will have 10-12 casualties and the total number of allowed triage reports for each run will be the number of casualties + 5.

Multiple triage reports will be permitted for a single casualty, and it is expected that the DARPA Command Post will award the best possible score out of the multiple triage reports. However, each triage report will count against the total number of allowed triage reports.

9.5.5 Final Ranking

For the Primary Triage Systems Competition, the final ranking will be determined based on the average of the top two scores of the team's three scored runs. In the event that multiple teams have an identical score, the team with the earlier non-zero scoring casualty report will be ranked higher.

10 Primary Triage: Virtual Competition Rules

10.1 Virtual Testbed

DARPA is investing in the development of the DTC Virtual Testbed, illustrated graphically in Figure 4, comprising of (1) DTC Simulator, an extensible Unreal Engine (UE5)-based simulation environment employing MetaHumans augmented with physiology data (2) automated testing and assessment tools; and (3) associated software support infrastructure. This suite of simulation tools is intended to support teams in both the Systems and Virtual Competitions as they develop and evaluate their approaches.

An initial version of the DTC Virtual Testbed and associated simulated sensing videos will be released in November 2023. This will be followed by regular releases with updates and enhancements (see Table 7). DARPA intends to continue adding significant improvements and new capabilities to the Testbed over the life of the challenge. Proposed capabilities for the Phase 1 challenge include support for user-generated paths, at least four sensing modalities (RGB camera with audio, thermal camera with audio, LiDAR), and at least 60 simultaneous casualties consisting of at least 4 injury classes (hemorrhage, burn, respiratory distress, and TBI) and at least 20 different human 3D models. The human models will vary by age (16 and older), size (varying BMI), and skin tone (using the 6 types from the Fitzpatrick scale).

Version	Release Date	Changes From Previous Version					Expected Use
		Nav Path	Injury Types	Sensing Modalities	# Sensor Options	# UGV / UAV Options	
0.1 (Phase 0)	11/23	Fixed	Hemorrhage, Burn, Respiratory Distress	RGB Thermal Audio	2	1 UGV 1 UAV	Phase 1 training / dev
0.3 (Phase 1) – Minor rel	3/24	Manual teleportation	Blunt Force Trauma		4	3 UGV 3 UAV	Phase 1 training / dev Phase 1 Workshop Qualification
0.5 (Phase 1) – Workshop	6/24	User Defined Waypoints	Chemical (respiratory depressant)	LiDAR	6		Phase 1 Workshop Phase 1 Challenge Qualification
1.0 (Phase 1) – Challenge	9/24 (DARPA)						Phase 1 Challenge
	11/24 (all)						Phase 2 training / dev Phase 2 Workshop Qualification
1.3 (Phase 2) – Workshop	3/25	Autonomy Interface	Bone fracture Amputation	RADAR Chem Detector	8	4 UAV	Phase 2 Workshop
1.5 (Phase 2) – Minor rel	6/25		Chemical (choking agent)		10	4 UGV	Phase 2 training / dev Phase 2 Challenge Qualification
2.0 (Phase 2) – Challenge	9/25						Phase 2 Challenge
	11/25						Phase 3 training / dev Phase 2 Workshop Qualification
2.3 (Phase 3) – Workshop	3/26		Polytrauma	Feedback Solicited	15	5 UAV	Phase 3 Workshop
2.5 (Phase 3) – Minor rel	6/26		Chemical (AChE inhibitor)			5 UGV	Phase 3 training / dev Phase 3 Challenge Qualification
3.0 (Phase 3) Challenge	9/26						Phase 3 Challenge
	11/26						Future Development

Table 7- DTC Virtual Testbed Anticipated Release Plan and Capability Growth

The DTC Virtual Testbed will use data from a physiology engine combined with injury timelines to simulate realistic MCIs and UxVs to provide a platform for training primary triage platforms and algorithms (Figure 4). The Testbed will also be used for challenge scoring at the end of each competition. A Virtual Testbed ICD that defines how teams provide input to the testbed (e.g., navigation waypoints / path, platform configuration, sensor configurations, etc.) and the required format / content of the casualty reports / logs will be published in early 2024.

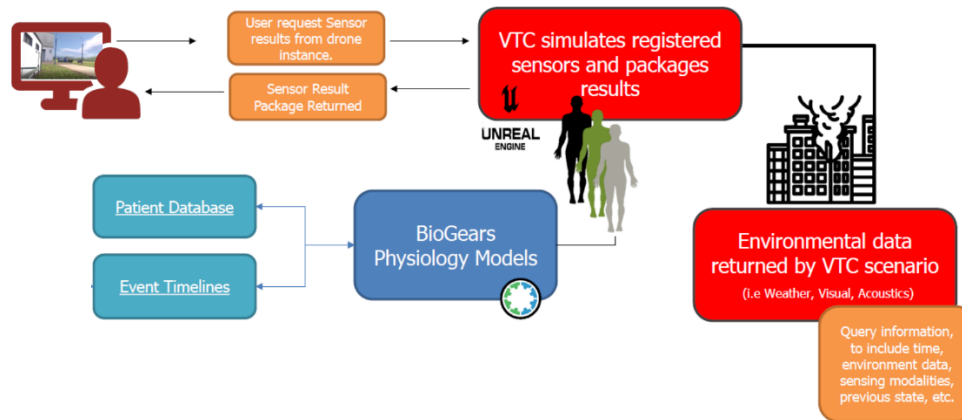


Figure 4 Virtual Testbed

10.2 Virtual - Technical Challenge Elements

The challenge elements used in the Virtual Competition will correspond to those of the Systems Competition to the fullest extent possible.

The Challenge competition courses will be designed to assess performance across various challenge elements, including: Degraded sensing, Obscuring obstacles, Terrain Obstacles, Dynamic Obstacles, and Dynamic Casualties. The challenge elements are expected to become progressively more difficult from Phase 1 to Phase 3.

1. *Degraded Sensing*: The courses are expected to include elements that range from constrained passages to large open fields, lighted areas to complete darkness, and clear to dusty conditions. Sensors will need to have the dynamic range to reliably operate in these environments. Dust, fog, smoke, talking, flashing light, heat spots, and loud background audio effects are within scope of this challenge element. Extreme temperatures, fire, and hazardous materials are not expected to be within scope.
2. *Obscuring obstacles*: Casualties may be fully visible to partially obscured to completely obscured, such as buried under a shallow layer of rubble. Sensor modalities capable of penetrating rubble will have an advantage in such situations. Stand-off sensor access to skin may be possible but cannot be assumed. Casualties may also be grouped with limbs, torso, or other body parts overlapping.
3. *Terrain Obstacles*: Systems will be required to demonstrate robustness in navigating a range of terrain features and obstacles. Terrain elements and obstacles may include constrained passages, large drops/climbs, inclines, and rubble. UAVs may be subject to atmospheric turbulence. The environments may include natural or human-made materials; structured or unstructured clutter; and intact or collapsed structures and debris.
4. *Dynamic Obstacles*: Responders, “walking wounded”, or other physical changes to the environment will test the agility of the system autonomy to identify and assess casualties.

5. *Dynamic Casualties*: Some treatable injuries may rapidly be fatal, so taking too long to find and assess casualties may result in the window for effective LSI to be missed. While competitors are not expected to re-evaluate casualties for changes in status, casualties who are not evaluated within an appropriate timescale may have a change in status (for example, progression of untreated hemorrhage or airway injury).

10.2.1 Challenge 1 Event Course Challenge Elements

The scale and complexity of competition courses is expected to increase each year. The design of the first Challenge Event courses is intended to assess the ability of teams to address the variety of challenging environments presented in a post-battle environment. In Year 1, platforms will travel through an area based on user defined waypoints and associated travel times / speed / dwell times. For competition one we will use challenge elements 2. Obscuring Obstacles and 5. Dynamic Casualties.

10.3 Virtual - Scored event submissions

10.3.1 Versions and releases

ARA, the DARPA performer providing the Virtual Testbed, will release four major versions; version 0.1 will be released in November 2023 (Phase 0), version 1.0 will be released in Fall 2024 (Phase 1), version 2.0 will be released in Fall 2025 (Phase 2), and version 3.0 will be released in Fall 2026 (Phase 3). ARA will also release at least one version between each of these versions. These “between” versions will be released prior to the days scheduled for virtual workshops for each challenge year.

The virtual competition scoring runs are expected to be performed by the IV&V teams (i.e., NOT self-evaluated). The challenge teams will not be permitted to witness the evaluations live.

10.3.2 Scored competition Scenarios

It is expected that teams will submit their solutions in the form of Docker Images to the DTC Virtual Portal where simulations will be run against unreleased competition scenarios that have not been seen by the challenge teams. It is anticipated that multiple scenarios will be devised for each simulated environment (e.g., all scenarios are the same incident and same virtual location, but will vary by the location of casualties, human model assigned to a casualty, and injury assigned to each human model). For Phase 1, it is expected there will be 3 different casualty scenarios that all use the same incident and virtual location. The competition scenarios, run scores, and logs will not be released until the event results are announced. Following the completion of a challenge, the event scenarios will be made available to competitors for training.

Each qualified team must submit a single solution per round to be scored. The submitted solutions will be evaluated against m number of competition scenarios to test the versatility of the solutions. Each competition scenario will, in turn, be evaluated over n replications (reps) to account for random variability. See

Competitors will be scored based on their ability to correctly evaluate casualties in simulation. Casualties will be presented at different severities with both physical indicators of injury and severity (i.e. lacerations and burns) as well as underlying physiological manifestations of injury (i.e. changes in respiration and heart rate). Casualties will be scannable in multiple sensing modalities, beginning with RGB, thermal, and audio at a minimum for Phase 1.

Distribution Statement ‘A’ (Approved for Public Release, Distribution Unlimited)

for a graphical depiction. The Event Score of the $m \times n$ runs is given by:

$$Event\ Score = \frac{1}{m} \sum_{i=1}^m \left(\frac{\sum_{j=1}^n run\ score\ e_{ij}}{n} \right)$$




Figure 5 - Virtual Competition approach to scoring competition scenarios

Competitors will be scored based on their ability to correctly evaluate casualties in simulation. Casualties will be presented at different severities with both physical indicators of injury and severity (i.e. lacerations and burns) as well as underlying physiological manifestations of injury (i.e. changes in respiration and heart rate). Casualties will be scannable in multiple sensing modalities, beginning with RGB, thermal, and audio at a minimum for Phase 1.

10.3.3 Solutions Submissions

For scored event submissions, it is expected that teams will submit their solution in the form of Docker Images to the DTC Virtual Portal, where it will be evaluated against the competition scenarios. Submissions must be self-contained and evaluated through an automated process. Entries that require additional user input or external commands will not be scored.

The solution submission window for the first challenge will open approximately 1 month prior to the challenge event. Each qualified team must submit a single solution to be scored. The submissions will be evaluated and the final results will be announced alongside the Systems Competition results in Fall 2024.

Challenge Event	Submission Window	Results Release
Challenge 1	Summer 2024	Fall 2024
Challenge 2	Summer 2025	Fall 2025
Challenge 3	Summer 2026	Fall 2026

Table 8 - Submission window for the Virtual track competitors

10.3.4 Human Supervisor

The submitted solutions will be evaluated with no external operator interfaces such as command line inputs or user interventions. Virtual Teams are required to develop self-contained solutions that navigate, search, and evaluate entirely autonomously without Human Supervisor interactions.

10.3.5 Staging Area

Each scored scenario begins with a Staging Area similar to the Systems Competition, inside which all platforms in the team may spawn. Platforms will not be spawned outside of the Staging Area. At spawn time, platforms are provided their ground truth location and orientation. Orientation along with a local coordinate system will be used for casualty reporting and navigation.

10.3.6 DTC Simulator Mechanics

The DTC Simulator utilizes a number of simulator mechanics to address the environmental accuracy of the competition and reduce the operational disparity between the Systems and Virtual Competitions. To that end, team submissions must fully utilize and not seek to circumvent use of simulator mechanics.

10.3.7 Run Termination

A scored run terminates upon any of the following conditions:

- Time Expiration: The given time expires before another termination criterion is met
- Completed run: All casualties have been successfully reported and assessed
- Casualty Reports Limit: The team reaches the maximum limit of allowed casualty reports, defined as $n + m$, where n is the number of casualties and m is 5, only the last one submitted will be used for scoring.
- Violation of minimum standoff distance. The user defined path for a robot breaches the minimum standoff distance based on total robot weight of the robot and all attached payload.
 - Less than or equal 250g: 1 meter.
 - More than 250g and less than or equal 5kg: 3 meters.
 - More 5 kg: 5 meters
- Violation of maximum height above terrain for UAV. The user defined path causes a UAV to exceed 30 meters above the terrain surface.

10.3.8 Score Disputes

Score Disputes are intended to provide teams a mechanism to submit a formal dispute or request for review by the Chief Judge. All score disputes must be sent by email to the DARPA Triage Challenge email address (triagechallenge@darpa.mil) within 48 hours of receiving competition log files. All disputes or requests will be reviewed by the Chief Judge in a timely manner. All decisions made by the Chief Judge are final.

10.4 Virtual - Preliminary Event Scenarios

10.4.1 Challenge Event Competition Environments

The scale and complexity of the environments is expected to vary across competition scenarios and across events. DARPA intends to release practice scenarios (in workshops) in advance of the Challenge Events to provide representative environments in which to develop and evaluate solutions.

In the year 1 competition events it is not expected that there will be space constrained portions of the course. In future years, it is possible that some portions of the environments will only be accessible via passages that are approximately one meter in height and/or one meter in width. The constrained spaces will be large enough for the available platforms to traverse but will limit visibility. The most constrained portions of the competition environments are not expected to be immediately at the entrance.

10.4.2 Preliminary Event Casualties

DARPA will announce the anticipated casualty types in advance of each competition event. The casualty types used in the Virtual Competition will be similar to those used in the Systems Competition. Casualty placement throughout the competition scenarios will be consistent across all runs for any given scenario but will vary across scenarios.

10.4.3 Challenge Event Run Duration

Each run will be between 15 and 30 minutes of in-simulation time.

10.4.4 Team Configuration

For Phase 1, the Team Configuration will be limited to one robot platform. In subsequent phases, the Team Configuration will be limited by the same constraints as the Systems Competition (i.e., maximum of 5 deployed autonomous systems). The DTC Virtual Testbed is expected to include a repository of Robot Operating System (ROS)-available mobile robot models and sensor payload models that will be available for teams to compose their Team Configuration.

10.5 Virtual - Communications and Score Reporting

10.5.1 Reporting casualties

Similar to the Systems Competition, teams will need to present relevant information (e.g., casualty health state report, casualty location, corresponding timestamps, etc.) to a virtual Base Station to provide near-real-time situational awareness updates and reports that are scored in a similar manner as the Systems Competition. To report a casualty for scoring, the casualty report must originate from a platform, be sent to the virtual Base Station, and include the location of the casualty being reported. Any casualty reports that do not include a location will count toward the allowed total, but will automatically receive a score of 0. Casualty reports that do report a location will be compared against the ground truth data to determine which actual casualty is closest to the reported location. The score will then be determined by comparing the reported casualty health state with the ground truth data at the time sensing was performed. It is expected that all scenarios in the first event will have at least 30 casualties and the total number of allowed scored reports for each run will be the number of casualties + 5

10.5.2 Log Files

At the termination of a run, relevant log files are generated. The log files include all triage reports (including the casualty's location), their corresponding timestamps (time sensing started, time sending ended, time assessment calculated), score updates, and other details. Additionally, log files allow replaying and viewing of the run by Competition Staff to ensure fair and consistent team performance in the virtual scenario in keeping with the rules and spirit of the DTC.

Distribution Statement 'A' (Approved for Public Release, Distribution Unlimited)

10.6 Virtual - Scoring Criteria

The goal of the DARPA Triage Challenge is to develop scalable, timely and accurate capture of novel injury signatures to enhance triage decision-making in austere, complex, and mass-casualty settings.

In the Primary Triage challenge (both Systems and Virtual), teams are evaluated based on accuracy and speed in assessing casualty condition using one or more autonomous platforms and stand-off sensors. Casualties are distributed throughout the competition course in a manner which rewards teams that are able to rapidly find and assess casualties. The location, nature of the casualties is not known in advance of a run by competitors and may be varied from run to run. Results for the Virtual competition will be announced at the prize ceremony on the last day of the competition event.

Upon identifying a casualty status, the deployed system must report its injury diagnosis and relevant clinical information to the evaluation system for scoring. In the Virtual competition, reports will be submitted to the virtual evaluation system. The detailed report format, protocol, and example implementation will be specified in the ICD. In both competitions, the reported casualty conditions are compared against coincident ground truth data for scoring.

10.6.1 Casualty Localization and Identification

For the phase 1 challenge, platforms are provided with a ground truth location and orientation. Identification and scoring will require localization of casualties relative to this origin position. Details will be released in future versions of this document.

10.6.2 Casualty Reports

Reports capture information about clinical condition of casualties relevant toward triage decision making. The casualty report is comprised of two parts: the first part contains vital signs and time-critical indicators of immediate need, and the second part contains more detailed assessment of injury condition and severity.

In future years, reports will include casualty location and may include additional clinical assessment features. The preliminary list of clinical features in the report relevant for scoring is shown in Table 9. Complete details on report contents and format will be found in the ICD.

10.6.3 Report scoring

A valid casualty report earns the team up to 10 points, with up to 5 additional bonus points for early reporting of time-critical information. Points are determined by comparing report contents to ground truth data from the same time. Categorical fields in the report are awarded points based on whether they match ground truth (GT); numerical fields in the report are awarded points based on whether they are within a predetermined range of the ground truth value.

To incentive fast assessment of time-critical information indicating immediate need of medical care, bonus points may be awarded for early casualty reports containing the following fields: Severe Hemorrhage, Respiratory Distress, and vital signs (Heart Rate and Respiratory Rate). Bonus points will be awarded for valid reports received by the DARPA Command Post within an initial “golden window” during a scored run. The duration of the “golden window” may be provided to teams prior to the challenge. Note that the report contents with bonus potential can be reported separately from those that do not have bonus potential, as described in the previous section.

Preliminary scoring criteria and bonus potential for clinical assessment in the casualty report are shown in Table 9. Casualty report definitions and examples will be found in the ICD.

Field	Values	Scoring Criteria
Severe Hemorrhage ^{1,3}	[Present, Absent]	2 if match ground truth (GT) 0 otherwise
Respiratory Distress ^{1,3}	[Present, Absent]	2 if match GT 0 otherwise
Heart Rate ^{1,4}	Integer	1 if within <i>n</i> of GT 0 otherwise
Respiratory Rate ^{1,4}	Integer	1 if within <i>m</i> of GT 0 otherwise
Trauma ²	Head: [Wound, Normal] Torso: [Wound, Normal] Upper Ext.: [Wound, Amputation, Normal] Lower Ext.: [Wound, Amputation, Normal]	2 if all match GT 1 if at least two match GT 0 otherwise
Alertness ²	Ocular: [Open, Closed, Not Testable (NT)] Verbal: [Normal, Abnormal, Absent, NT] Motor: [Normal, Abnormal, Absent, NT]	2 if all match GT 1 if at least two match GT 0 otherwise

Table 9 - Preliminary casualty report clinical assessment with scoring criteria.

¹ Time-critical features comprising the first part of the report.

² Detailed injury and severity assessment comprising the second part of the report.

³ Correct response receives 2x bonus if reported within “golden window”.

⁴ Correct response receives 1.5x bonus if reported within “golden window”.

10.6.4 Time

Time is measured from the start of the run and scoring will be based off of two timestamps: the assessment timestamp and the report timestamp. The assessment timestamp is provided by the deployed system in the report and indicates the time at which the casualty’s condition was assessed. This timestamp is used to determine the appropriate ground truth for scoring. The report timestamp is the time when the valid report is received at the DARPA command Post, and it determines any bonus points awarded.

10.6.5 Report limits

To discourage guessing and preserve system bandwidth, the DARPA Command Post will limit the total number of scored reports. Any further reports are rejected. In phase 1 it is expected that each run will have 10-12 casualties and the total number of allowed triage reports for each run will be the number of casualties + 5.

Multiple triage reports will be permitted for a single casualty, and it is expected that the DARPA Command Post will award the best possible score out of the multiple triage reports. However, each triage report will count against the total number of allowed triage reports.

In the Primary Triage Virtual Competition year 1, it is expected that all scenarios have at least 30 casualties and the total number of allowed scored reports for each run will be $N + 5$. For this Competition, exact duplicate reports will not be counted against the total number of reports.

10.6.6 Final Ranking

For the Primary Triage Virtual Competition, the final ranking will be determined based on each team's event score as described in Section 10.3.2. In the event that multiple teams have an identical event score, those teams will be ranked in ascending order according to the report timestamp of the latest non-zero scoring casualty report.

11 Secondary Triage: Data Competition Rules

11.1 Data - Illustrative Scenario

The objective of the Secondary triage competition is to identify physiological signatures of injury derived from data captured by non-invasive sensors (contact-based or stand-off) to enable anticipatory decisions and prioritization for medical evacuation and care. Performers will develop algorithms that detect signatures in these data streams to provide decision support appropriate for austere and complex pre-hospital settings. Of particular interest are early signatures indicating need for LSIs against conditions that medics are trained and equipped to treat during secondary triage, such as hemorrhage and airway injuries.

The Secondary triage competition is virtual only, and will use DARPA provided de-identified, multi-modal physiological data from trauma patients in diverse settings and cohorts provided by the DARPA **Research Infrastructure for Trauma with Medical Observations (RITMO)** program. DARPA will provide access to a subset of these data for algorithm training, and evaluate competitor algorithms on holdout test data in end-of-phase challenge events. Challenge events will become progressively more complex and realistic from Phases 1 to 3.

11.2 Data - Technical Challenge Elements

The Challenge events will be designed to assess performance across various challenge elements, including: multiple data sources, multiple data inputs, raw data, and degraded data. The challenge elements are expected to become progressively more difficult from Phase 1 to Phase 3.

1. *Multiple data sources*: With varied source populations, patient demographics, types of injury, and standard clinical operating procedures. Each data set may have a different standard set of sensor readings, with different added sensors in each phase of DTC. Approaches must demonstrate robustness across a range of settings.
2. *Multiple data inputs*: Potentially including static data (e.g., mechanism of injury or anatomical injury pattern); multiple simultaneous, continuous streams of high-frequency waveforms; and point-of-care imaging data.
3. *Raw data*: Data as it comes from the sensors and health record systems (i.e., without any cleaning), with the noise, aberrant values, and dropouts that occur in clinical environments.
4. *Degraded data (year 2 and 3)*: DARPA also may inject additional challenges that can be expected in battlefield and civilian pre-hospital settings, such as severe degradation or total loss of a particular sensor, to test the robustness of competitor strategies to such plausible scenarios.

11.3 Data Types

Competitors in the Data Challenge will be given data specifications for the modeling/training data that will be used. DARPA-provided datasets originate from the DARPA Research Infrastructure for Trauma with Medical Observations (RITMO) program. The dataset contains at least 3000 cases of pre-hospital or hospital data for year 1 and may include: discrete data (text or numeric values): patient characteristics, outcomes, procedures, labs; Continuous data: (electrocardiogram) ECG, photoplethysmography (PPG), respiration rate (RR), heart rate (HR), SPO2, blood pressure (BP), Arterial-line BP. Data dictionaries will be provided with details on data format and proper interpretation.

While all data sources will be made available for training, only a subset of the data will be available at test time for predicting LSIs. These will include data sources that would feasibly be available to medics *in situ* during treatment, excluding information about future events (e.g., outcomes) or contextual information typically collected post-treatment (e.g., medical history).

11.4 Data - Scored event submissions

11.4.1 Solutions Submissions

For scored event submissions, teams must submit their solution in the form of Docker images (saved as a .tar file) to the Challenge Submission folder in AWS, where it will be evaluated against a held-out RITMO testing dataset with predetermined compute resources. Solutions must be self-contained within the Docker image and must be submitted to an *S3 Submission Bucket* which will be automatically processed through an evaluation system. Guidelines on required Docker image configuration will be included in the ICD.

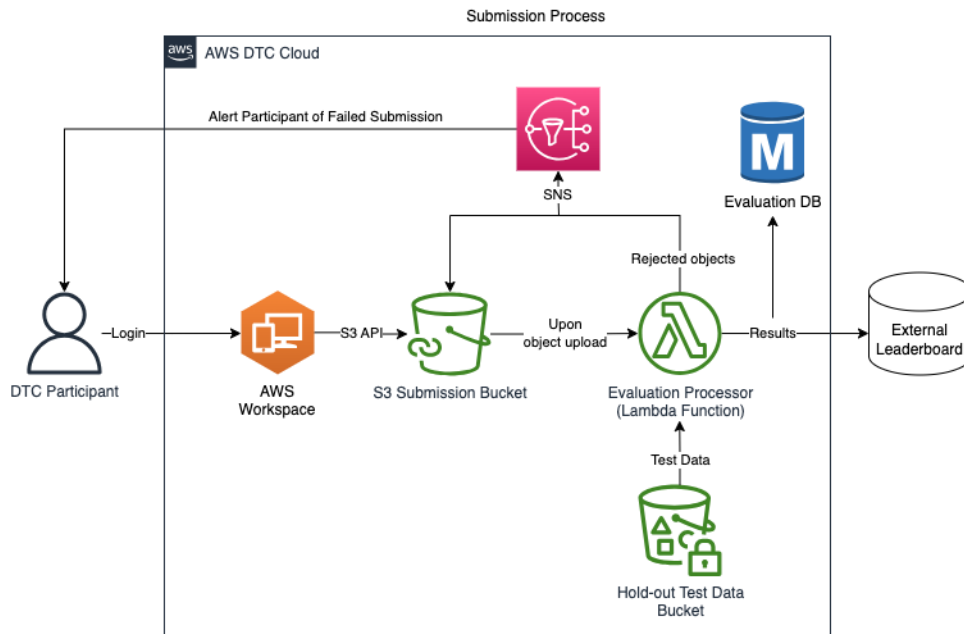


Figure 6 - Data Challenge AWS Submission Process. Participants will send their Docker solution to an S3 Submission Bucket within their team's workspace using the S3 API. Once the transfer is completed, an automatic process will evaluate the solution against a hold-out dataset. Results of the solutions will be published to an external, public leaderboard. The submitting team will receive an email notification regarding the status of the evaluation results

The solution submission window for the first challenge will open approximately 2 months prior to the challenge event. Each qualified team must submit a single solution to be scored. The submissions will be evaluated and the final results will be announced alongside the Systems and Virtual Competition results at the Competition Awards ceremony.

Challenge Event	Submission Window	Results Release
<i>Challenge 1</i>	<i>Summer 2024</i>	<i>Fall 2024</i>
<i>Challenge 2</i>	<i>Summer 2025</i>	<i>Fall 2025</i>
<i>Challenge 3</i>	<i>Summer 2026</i>	<i>Fall 2026</i>

Table 10 - Submission window for the Data Competition teams.

11.4.1 Submission scoring

Each qualified team must submit a single solution per round to be scored. Team solutions will be evaluated under a simulated online prediction paradigm using m held-out test cases. Submitted solutions will be given incremental data within consecutive time intervals to approximate streaming source data, and any predicted LSIs over time will be compared against ground truth. Solutions may predict each LSI at most once per case.

Scores will reflect relative performance of each teams' solution compared to the other solutions. Within each case, teams will be ranked according to the accuracy and timeliness of their responses. Points will then be awarded to the top-ranking teams. The Event Score is the sum of points received across all cases. Details on prediction scoring and team ranking are in Section 12 of this document.

$$\text{Event Score}_k = \sum_{i=1}^m S_{i,k}$$

Figure 7- Data Competition approach to scoring team solutions against each other.

The final competition scores and any supporting materials will not be released until the event results are announced.

11.4.2 Human in the loop

The submitted solutions will be evaluated with no external operator interfaces such as command line inputs or user interventions. Teams are required to develop self-contained solutions that predict LSIs entirely autonomously without Human Supervisor interaction. Entries that require any user input or external commands will not be scored.

11.4.3 Run Termination

Teams will be scored against each other within each case in the test dataset. A scored run for a given case terminates upon any of the following conditions:

- Time Expiration: The end of the evaluation window was reached or the evaluated solution processing time exceeded the time limit, whichever occurs first
- Prediction Saturation: The evaluated solution predicted the need for *all* LSI types
- Patient Expiration: The given patient expired before another termination criterion was met

11.4.4 Score Disputes

Score Disputes are intended to provide teams a mechanism to submit a formal dispute or request for review by the Chief Judge. All score disputes should be sent by email to the DARPA Triage Challenge email address (triagechallenge@darpa.mil) within 48 hours of receiving competition log files. All disputes or requests will be reviewed by the Chief Judge in a timely manner. All decisions made by the Chief Judge are final.

11.5 Data - Preliminary Event Scenarios

11.5.1 Preliminary Event Competition Environments

All computing operations and analysis involving the RITMO dataset must occur within the AWS ecosystem throughout the duration of the challenge. Each team will be provided an isolated, network-restricted AWS Workspace domain which will allow teams to manage storage and computing resources. Team members will be provided Workspace user accounts and can instantiate computing resources (AWS EC2, SageMaker, etc.) on demand within these Workspaces. To use GPU resources, teams would deploy AWS SageMaker instances. Each team is allowed a maximum of running three SageMaker instances at a time. The table below shows the available resources, their compute specification, the maximum number of running instances, and the price per hour:

Instance Type	Compute Specifications				Purpose	Limit (per team)	Price per Hour (\$)
	Number of CPU	RAM (GB)	Number of GPUs	GPU Memory			
ml.t3.medium	2	4	0	0	General Purpose	3	0.05
ml.m5.large	2	8	0	0	General Purpose	3	0.115
ml.m5.2xlarge	8	32	0	0	General Purpose	3	0.461
ml.c5.xlarge	4	8	0	0	Compute Optimized	3	0.204
ml.c5.2xlarge	8	16	0	0	Compute Optimized	3	0.408
ml.c5.9xlarge	36	72	0	0	Compute Optimized	1	1.836
ml.r5.8xlarge	32	256	0	0	Memory Optimized	1	2.419
ml.r5.16xlarge	64	512	0	0	Memory Optimized	1	4.838
ml.p3.2xlarge	8	61	1	16	GPU – General	1	3.825
ml.p3.8xlarge	32	244	4	64	GPU – General	1	14.688
ml.g4dn.8xlarge	32	128	1	16	GPU – Training	1	2.72
ml.g5.4xlarge	16	64	1	24	GPU – Inference	1	2.03
ml.g4dn.xlarge	4	16	1	16	Accelerated Computing	3	0.7364
ml.g4dn.2xlarge	8	32	1	16	Accelerated Computing	3	0.94
ml.g4dn.8xlarge	32	128	1	16	Accelerated Computing	3	2.72

Table 11 - AWS processor options.

Teams will be allocated a fixed budget at the beginning of each phase to spend on storage and computing resources. It will be the responsibility of the team to manage their budget expenditure and resource consumption throughout the competition. All teams will receive a weekly email notification regarding the status of the team's expenditure and resource usage, and email notification when teams reach specific budget drawdown levels. Any remaining funds (up to 25%) at the end of a phase may be rolled into the

subsequent phase. The Government's obligation for AWS funding under the DARPA DARPA Triage Challenge is subject to the availability of appropriated funds.

While working within Workspaces, participants are allowed to download data from the internet, allowing them to browse external websites, and download opensource packages. However, participants are restricted from uploading data to ensure data does not leak outside the AWS ecosystem.

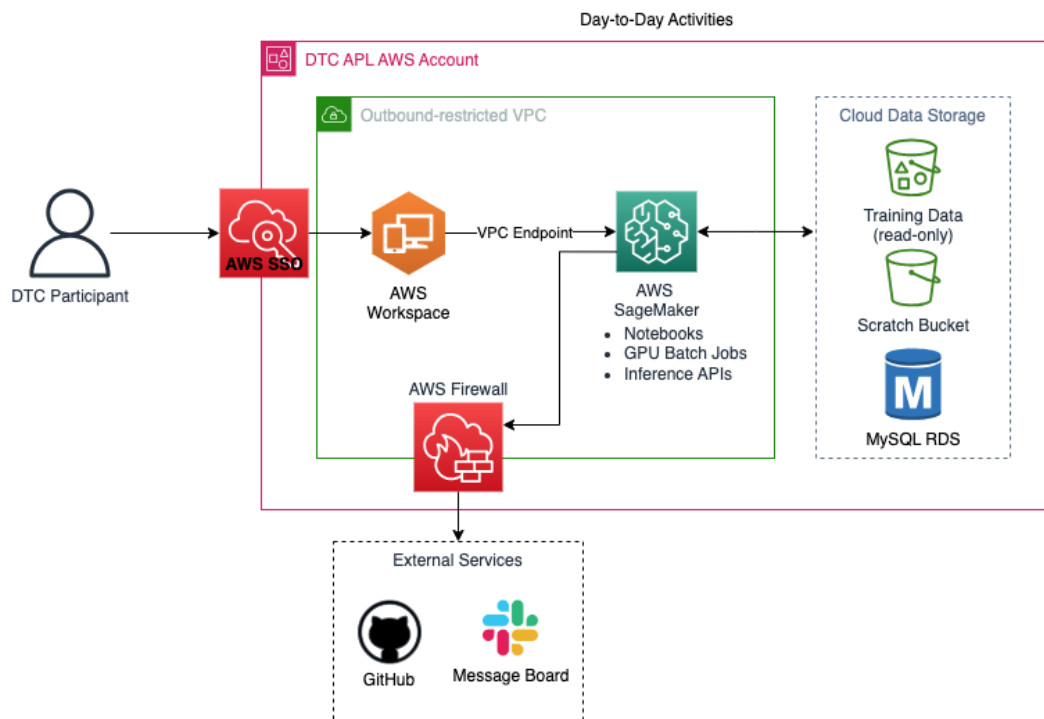


Figure 7 - DTC Participant AWS Architecture Diagram. Participants will perform all computations pertaining to the RITMO dataset within the provided AWS ecosystem. Each participant will login to their designated AWS Workspace which allows them to access the RITMO dataset stored in either a MySQL Database or an S3 bucket. Participants will be able to instantiate computing resources (e.g., AWS EC2, AWS SageMaker, etc.) within the Workspace. Each team will also be provided a Scratch Bucket to share data or files among the team members.

11.5.2 Preliminary Event LSIs

As part of the RITMO dataset, treatments and related clinical actions have been identified and grouped into LSI categories based on shared injury patterns and treatment paradigms. Teams in the Data competition are tasked with predicting the occurrence of these LSI categories, not the specific treatment or clinical action. For brevity, the LSI categories will be referred to as LSIs for the remainder of this document. LSIs with timestamps will be provided alongside the training and test datasets.

The preliminary list of LSIs includes: airway & respiration, vascular access & monitoring, bleeding control, vasoactive & cardioactive medications, crystalloid products, blood products, cardiovascular procedures, thoracic & other procedures, chest decompression, neurological products & procedures, and RSI sedation medications. These LSIs may change in future years to reflect new data sources or changing priorities of the challenge. Details on the LSIs can be found in the data dictionary provided alongside the dataset.

11.5.3 Prediction Task

Submitted solutions will be evaluated at predefined evaluation timepoints along each case, ingesting incremental continuous and discrete data within consecutive windows preceding each timepoint. At each timepoint, the task is to predict the set of LSI types present within a time window in the future relative to the observation time-point. See Section 11.6.3 for details on the prediction task scoring criteria.

11.5.4 Preliminary Event Run Duration

The event run duration is defined as the cumulative processing time for solutions to produce prediction responses across all evaluation timepoints across all test cases. DARPA will set an evaluation duration limit based on the number of time points at a later date.

11.6 Data - Scoring Criteria

In the Secondary Triage (Data) Challenge, teams are evaluated based on sensitivity, specificity, and speed in predicting future LSIs from physiological signals and contextual health information. Results for the Data Competition will be announced at the prize ceremony on the last day of the competition event.

Submitted solutions in the data competition will be evaluated for accurate and early prediction of the need for LSIs using a set of held-out test cases. To simulate online prediction, solutions will be tested at multiple evaluation timepoints along each case. At each evaluation timepoint, solutions will be provided with incremental data from the past and any predicted LSIs will be compared against the ground truth set of LSIs present in the future, where past and future are defined relative to the evaluation timepoint. Details on the input and output formats for data and LSI predictions, respectively, will be found in the ICD.

The Event Score is based on the performance of each team relative to the other teams. Within each case, teams will be ranked according to the accuracy and timeliness of their LSI predictions. Top-ranking teams will then receive a point for out-performing the other teams. The sum of points across cases in the test set will determine the winning team.

11.6.1 Definitions

All solutions will be evaluated at a predetermined set of *evaluation timepoints* within each *case*, where a case encompasses the pre- and in-hospital data and timestamped LSIs from a single hospital admission. At each evaluation timepoint, timestamped data within the *observation window* will be provided to submitted solutions for the prediction task, where the observation window begins with the start of care and ends at the evaluation timepoint. The *prediction target* is the unique set of ground-truth LSIs within the *prediction window*, the time window beginning 15 minutes after the evaluation timepoint and extending to the end of care or four hours after hospital admission, whichever comes first.

The gap between observation and prediction windows represents the minimum lead-time necessary for actionable clinical predictions that would provide value for medical resource allocation and planning. Any LSIs occurring outside of the prediction window (i.e., within the observation window or the intervening gap) are excluded from the prediction target. See Figure 8 for an illustration of the evaluation timepoint and surrounding windows, and how the prediction target is derived from the underlying LSI instances.

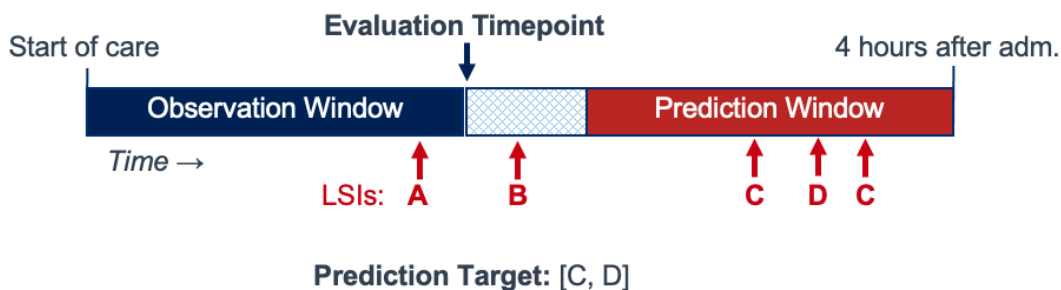


Figure 8 - Illustration of evaluation at a single timepoint. Solutions generate predictions using data occurring within the observation window preceding the evaluation timepoint. The prediction target is the set of timestamped LSIs (C and D) present in the prediction window, which begins after the evaluation timepoint and a predefined gap. LSIs occurring before the prediction window, whether within the observation window (A) or the intervening gap (B), are not included in the prediction target. Note that the prediction target is the set of unique LSIs present in the prediction window, not the frequency or order of LSIs.

11.6.2 Online prediction

Solutions will be tested at one or more evaluation timepoints along each case. At each timepoint, solutions will be provided with incremental data building upon data provided at any earlier timepoints, and solutions will be given the opportunity to predict the set of LSIs present in the prediction window relative to the current evaluation timepoint. At test time, solutions will be provided with the timestamp of the evaluation timepoint, timestamped data, and a case identifier. Evaluation timepoints will be tested in temporal order, and after predictions from the final evaluation timepoint within a case have been collected, the sequence of predicted LSIs will then be scored against the sequence of ground truth LSIs over the case.

To incentivize efficient processing and ensure time-bounds on the evaluation, the run-time of solutions will be recorded for each prediction. There will be a predetermined maximum processing time, not to exceed the minimum lead-time, after which the prediction is assumed to be the empty set and the evaluation will continue to the next evaluation timepoint or the next case, whichever is appropriate. The processing delay of correct predictions in each case will factor into the scoring, as described in the next section.




Figure 9 - Example predictions over a single case. The solution from “Team Z” is offered an opportunity to predict LSIs at three evaluation timepoints along the case, indicated by index t and the windowed bars above. At each timepoint teams are provided with incremental data building on previous data provided at earlier timepoints (navy bars). Without loss of generality, LSIs are abstracted to the symbols {A, B, C, D, E}. Y_t is the set of ground truth LSIs in the prediction window at evaluation timepoint t , and \hat{Y}_t is the LSI predictions at evaluation timepoint t . Arrows indicate processing delay at each evaluation timepoint. Prediction score is based on the sequence of predictions \hat{Y}_t and the processing delay for correct predictions across the entire case.

11.6.3 Scoring predictions

For each case, teams will be ranked in descending order according to a time-sensitive similarity metric comparing LSI predictions to ground truth LSIs over time. Teams with the same similarity metric will then be ranked further in descending order according to a timeliness metric. The similarity metric and timeliness metric are described in the following sections.

After ranking by similarity metric and by cumulative prediction lead-time, the top k teams receive a point towards their event score. The winning team has the highest event score, or equivalently, the highest number of cases where they out-performed $N-k$ other teams, where N is the total number of teams. The number of top-ranking teams that receive a point within each case, k , will be selected such that there is a single highest scoring team.

11.6.4 Similarity metric: Intersection-over-Union

The time-sensitive similarity metric is based on the common intersection-over-union statistic: it rewards for correctly predicted LSIs and penalizes for false positive, missing, or late predictions. It is the ratio of count of correct LSI predictions across timepoints to the count of all predicted LSIs and all ground truth LSIs across the case. This similarity metric is bounded between 0 and 1, where 1 indicates correct prediction of all LSIs without any false positives and 0 indicates that none of the ground truth LSIs were correctly predicted in time.

$$\text{IoU} = \frac{|\bigcup_{t=0}^T [\hat{Y}_t \cap Y_t]|}{|\bigcup_{t=0}^T [\hat{Y}_t \cup Y_t]|}$$

Figure 10 - Time-sensitive similarity metric based on Intersection-over-Union (IoU). The denominator is the count of all LSI predictions \hat{Y}_t and all ground truth LSIs Y_t and the numerator is the count of all correct LSI predictions across the case and the denominator is the count of all LSI predictions \hat{Y}_T and all ground truth LSIs Y_0 across the case, where evaluation timepoints span $t=0$ to $t=T$, where $t=0$ and $t=T$. For cases with no LSIs, the similarity metric collapses to 0 for solutions that falsely predict any LSIs at any evaluation timepoint, and 1 for solutions that correctly do not predict any LSI at all evaluation timepoints. In this special case, any team that scores 1 receives a point toward their event score.

11.6.5 Timeliness metric: Prediction Lead-Time

The prediction lead-time measures how early a given solution correctly predicted LSIs within a case. It is the sum of time intervals between the simulated prediction timestamp and the ground truth LSI timestamp, where the *prediction timestamp* is the evaluation timepoint plus the processing time of the solution, as recorded by the evaluation system at run-time.

The prediction lead-time is accumulated across each correctly predicted LSI occurrence, thereby rewarding the solution with an earlier correct prediction of multiple LSI occurrences of the same type.

$$\text{PLT} = \sum_{\substack{\text{Earliest correct predictions} \\ \text{for each LSI type}}} \text{LSI timestamp} - \underbrace{(\text{Evaluation timepoint} + \text{processing delay})}_{\text{Simulated prediction time}}$$

Figure 10 - Timeliness metric, Prediction Lead Time (PLT). The time difference between the ground truth LSI timestamp and the simulated prediction time, accounting for processing delay, is summed across all correct predictions. For cases with correct prediction of multiple LSI instances from the same group, they count as separate elements in the sum. Using the example in Figure 10, there are three items in the sum: the time difference between both ground truth instances of [A] and the correctly predicted [A], and the time difference between the ground truth instance of [D] and the correctly predicted [D].

11.6.6 Final Ranking

For the Secondary Triage Data Competition, the final ranking will be determined based on each team’s event score. In the event that multiple teams have an identical event score, tiebreakers will be applied by reducing *k*, the number of top-ranking teams receiving a point for each test case, until a winner emerges. In the event that multiple teams have an identical event score with *k*=1, those teams will be ranked in descending order according to the sum of prediction lead-times across all test cases.

12 Appendix 1 DTC Glossary

Chief Official – Program manager or higher DARPA authority for the DARPA Triage Challenge.

Systems Competition – Primary Triage Competition run with actors on a real course (Track A, B).

Virtual Competition – Primary Triage Competition run on a virtual platform (Track C).

Data Competition – Secondary Triage Competition (Track D, E).

Base Station – One or more computers or controllers that serve as the interface between the systems, the DARPA Command Post, and the Human Supervisor.

Chief Judge – DARPA-designated individual with the sole and final authority to make any decisions related to the rules or scoring.

Competition Course – Physical or virtual environment in which deployed systems are expected to explore, and search for casualties.

Course Official – DARPA-designated individual that is based in each Staging Area to apply and enforce the rules and make safety-related decisions, with decision-making authority only superseded by the Judge and Chief Judge.

DARPA Command Post – Computer interface which receives casualty reports and map updates from teams and returns score updates. Also refers to the main headquarters where the DARPA staff execute the Challenge.

Human Supervisor – Team-designated individual permitted to interface with the Base Station, provide high-level interactions with the deployed systems, use wireless communications, and access both course data and status data.

Judge – DARPA-designated individual with authority to make decisions related to rules, scoring, and safety, with decision-making authority only superseded by the Chief Judge.

Pit Crew – Team personnel permitted in the Staging Area to assist with operations tasks such as physically deploying the systems, performing repairs, modifying software or firmware, and changing batteries.

Safety Officer – Team-designated members of the Pit Crew responsible for preserving the safety of personnel and property, activating emergency stop transmitters, and/or operate remote controls for safety purposes.

Staging Area – Specified area immediately outside of the Competition Course entrance from which teams deploy their system.

Starting Gate – Installed structure or existing entrance which serves as the threshold between the Staging Area and the Competition Course.

Starting Gate Fiducial – An easily identifiable object attached to or near the Starting Gate to assist teams to align with the official coordinate frame in which casualties are reported. These may include 2D barcodes, reflective targets, or survey prisms.

Team Garage Coordinator- DARPA-designated individual supporting team prep.

Team Lead – Team-designated individual responsible for making official team decisions (e.g., termination of a run) and communicating with the DARPA Competition Staff.