

Spectrum Collaboration Challenge

Phase 3 Frequently Asked Questions (FAQ)

October 4, 2019

Revision 9



Defense Advanced Research Projects Agency
675 North Randolph Street
Arlington, VA 22203



Revision Summary

Section	Revision	Description	Date
Q66	9	New	10/4/19
Q60-Q65	8	New	9/16/19
Q57-Q59	7	New	8/30/19
Q54	6	Update to Scrimmage 6 Submission Deadline	8/29/19
Q55 – Q56	5	New	8/12/19
Q50 – Q54 Q27	4	New Clarified, which flows have mandated outcomes	8/1/19
Q47 – Q49	3	New	5/10/19
Q41 – Q46 Q7	2	New Revised	4/15/19
Q1 – Q40	Initial	Carryovers from Phase 2	2/1/19

DARPA Spectrum Collaboration Challenge (SC2) Phase 2 Frequently Asked Questions

Q66: Per section 2.3 of the SCE Scoring Procedures, how will tie scores in the final round of SCE resulting from the points awarded per section 2.3 of the SCE Scoring Procedures be handled?

A66: The prize award amounts for all positions which result in a tie will be pooled and equitably distributed among the tied teams.

Q65: Will SCE logs be made available to competitors post-SCE?

A65: Yes, logs will be provided by November 8th 2019.

Q64: Per Q51, MCHM saturation is possible between CIRN nodes, however, is it guaranteed that the MCHM observer nodes will never experience saturation?

A64: It cannot be guaranteed that the observer nodes will *never* experience the saturation conditions described in Q51. However, brief and rare saturation of the observer nodes will have no impact to scoring.

Q63: Are the TOS fields used in the calculation of scores?

A63: No, TOS fields are superseded by mandate points-values.

Q62: When an environment update changes the value of *scoring_point_threshold* does it go into effect immediately?

A62: Yes.

Q61: Will there be matches where two instances of the same team play against each other?

A61: No scored matches of SCE will use multiple instances of a radio network submitted by the same team. DARPA may however collect data in this configuration for use in post competition presentations and publications.

Q60: Will official Play-in Round rankings be released?

A60: DARPA will release official Play-in Round rankings on September 27th after final competitor submissions for SCE are received.

Q59: What is the CIL requirement for reporting current score and Individual Mandate performance?

A59: When a team is below the Ensemble Threshold, the team must report IMs achieved and points earned accurately. When a team is above the Ensemble Threshold, there is no accuracy requirement except to state “above threshold”. However, teams may share additional information to improve final Match Score.

Q58: In the Scoring Procedures document, section 2.2.1 (“Elimination Rounds: Round-Robin”), under “Tiebreaker #1”: is “the total time above the Ensemble Threshold” a measure of individual team performance, or of collective ensemble performance?

A58: For the purpose of resolving ties in round-robin scores, “Above Threshold Time” for network N is defined by the total number of measurement periods in which network N’s Ensemble Score (RMPS of the lowest-performing network) met or exceeded the Ensemble Threshold for all matches in the round.

Q57: In the Scoring Procedures document, section 2.2.2 (“Elimination Rounds: Knockout”), under “Tiebreaker #1”: is “the total time above the Ensemble Threshold by each competing team” a measure of individual team performance, or of collective ensemble performance?

A57: For the purpose of resolving ties in knockout match scores, “Above Threshold Time” for network N is defined by the total number of measurement periods in which network N’s individual Relative Measurement Period Score (RMPS) met or exceeded the Ensemble Threshold during the knockout match.

Q56: What is the number of teams that will play in an SCE Final Round match?

A56: See Section 2.3 of the SC2 Championship Event Scoring Procedures.

Q55: Per Q6, are incumbents present in the throughput performance testing scenario?

A55: There will be no incumbents present in the throughput performance testing scenario.

Q54: Per Q40, how will teams be evaluated during the play-in round to gain entry to the SC2 Championship Event (SCE)?

A54: Teams must submit an “image” for Scrimmage #6 by 12:01 AM US Eastern Time ~~September 3rd 2019~~ **September 4th 2019** according to the submission logistics documented here:

<https://sc2colosseum.freshdesk.com/support/solutions/articles/22000235720-phase-3-scrimmage-logistics>

Submissions will be tested for compliance with the criteria laid out in Q6 and Q7 of this FAQ. Submissions which fail to meet these criteria will not be accepted, nor evaluated in the play-in round. The play-in round will exhaustively evaluate all 3-team combinations in the *Alley's of Austin* scenario. Submissions will be ranked according to their Round-Robin Score (as defined in the Section 2.2.1 of the SC2 Championship Event Scoring Procedures). The top 10 ranked teams will be admitted to the SCE.

Q53: What is the procedure for final submission for the SC2 Championship Event?

A53: Final submissions must be made according to the logistics posted here:

<https://sc2colosseum.freshdesk.com/support/solutions/articles/22000243206-sce-submission-logistics>

Q52: If the final image submitted by a team for the SC2 Championship Event (SCE) fails to pass qualification criteria laid out in Q6 and Q7, is that team eliminated from SCE?

A52: No, the failing submission will be replaced with the team's most recent submission which passed the qualification criteria.

Q51: Is it possible for competitors to overdrive the processing within the MCHM digital channel emulator in Colosseum?

A51: Yes. The MCHM's channel emulation processing is bounded by hardware constraints of both the radio channel RF front-ends and the digital signal processing back-end. All RF front-ends can be overdriven when competitor RX and TX gains are improperly configured. Additionally, mismanaged gains can also cause the digital signal processing back-end of MCHM to be overdriven. Overdriving MCHM's digital back-end causes arithmetic overflow, either 1) into the sign bit or 2) clamping to the maximum value. These typically manifest as either phase discontinuities or saturation. These nonlinearities, when observed in the frequency domain, appear as broadband noise. Competitors are able to prevent this by managing their analog and/or digital gain settings. Either fixed or dynamically adjusted gain settings may be used to avoid overdriving MCHM.

Q50: Per the scoring document, “Throughput is measured as bits per second of IP traffic delivered to the traffic generator sink at the destination node.” How are ip packet sizes counted?

A50: Scoring uses what MGEN reports as the message length, which is defined in their official documentation:

<https://downloads.pf.itd.navy.mil/docs/mgen/mgen.html# MGEN Message Payload>

Q49: Is the "Scenario-local team" index that is used in the score json files equal to the tens-digit of the RFNode_ID in the batch_input.json file?

A49: The scenario-local team index used in scoring JSON files corresponds to the order in which the teams are listed (by letter, from "A" to "E" in a 5-team match) in Colosseum match metadata files. It's important to note that if Colosseum ever assigns node numbers other than 1-10 to "Team A" (for example), the scoring engine will still label that as "scoring team 0". Also important to note that this is not the same as "the tens digit" because Colosseum node numbering begins at 1, not at 0.

Q48: Per Q44, How can I determine whether I am meeting the MO for a File burst flow?

A48: Measure the delay for each File traffic packet. If 90% of the File traffic packets offered in a MP are delivered with a delay less than the file transfer deadline, then the File burst flow MO is satisfied in that MP.

Q47: Will CIL compliance be verified outside of the qualification checks outlined in Q7?

A47: Per SC2 Rules Section 3.2.1, DARPA will conduct testing to verify the integrity of the competition. This testing includes the qualification tests prescribed by Q6 and Q7. Prior to the awarding of prizes, DARPA reserves the right to ensure all winners have maintained the competition's integrity. This includes (but is not limited to) checks to ensure there is no attempted or actual exploitation of the scoring engine, CIL compliance verification, nor purposeful or artificial limitation of performance that hinders an ensemble's ability to meet or exceed the Ensemble Threshold.

Q46: If an Individual Mandate's PTs are met for 10 consecutive MPs with no gate violations (and the SSP is 10 seconds), are points awarded in the 10th MP or the 11th MP?

A46: Points are awarded in the 10th MP.

Q45: Will the measurement periods in the final 15 seconds of a match, after the traffic generator has stopped sending traffic, be scored or unscored?

A45: Unscored.

Q44: What is the minimum time spacing between 2 subsequent bursts in a file transfer?

A44: There is no guaranteed minimum time.

Q43: For SCE, will the Ensemble Threshold (ET) ever be set to zero?

A43: No.

Q42: For SCE, what is the maximum IP packet size?

A42: 1500 bytes.

Q41: For SCE, will any scenarios require multi-hop routing?

A41: No SCE scenarios are designed with this objective in mind. However, the ability to "close" a radio link depends heavily on the specific radio design including parameters such as the modulation and error correction type, the baud rate used, and the bandwidth of any digital filters. As such, DARPA cannot guarantee that all radio links can be closed.

Q40: Will all teams admitted to Phase 3 participate in SCE?

A40: No, only 10 teams will be admitted into SCE, based on a ranked "play-in" round to take place coincident with Scrimmage 6.

Q39: How is a Measurement Period (MP) scored if the total data offered to a radio in the MP is less than the specified *min_throughput_bps* Performance Threshold (PT)?

A39: During such a MP, the *min_throughput_bps* PT is reduced to match the offered data.

Q38: Are Measurement Periods (MPs) defined by the integer portion of the transmit time recorded in the .drc logs?

A38: No. The start of the first 1-second-long measurement period is defined by the match start time reported by Colosseum. Packets are attributed to 1-second-long MPs according to their sent timestamp relative to the match start time.

Q37: In order to measure throughput performance per Q6, is the TRPR tool used?

A37: No. DARPA has developed custom tools which calculate throughput using logs produced by MGEN.

Q36: In order to measure throughput performance per Q6, how often is throughput calculated?

A36: Throughput is computed each Measurement Period (MP) as defined in the SCE Scoring Procedures document.

Q35: Is there a complete list of all scenario bandwidths a CIRN will be asked to support for SCE?

A35: A CIRN may need support all of the following scenario bandwidths for SCE: 5 MHz, 8 MHz, 10 MHz, 20 MHz, 25 MHz, 40 MHz.

Q34: For SCE, What is the maximum total throughput requested of a single node?

A34: 25 Mbps.

Q33: What is the maximum value for *max_latency_s*?

A33: There is no upper bound on *max_latency_s*.

Q32: Can file transfers overlap such that a new packet burst occurs before *file_transfer_deadline_s* has elapsed from the previous burst?

A32: Yes.

Q31: For SCE, what is the maximum number of flows for a node to handle simultaneously?

A31: There is no upper bound on the number of flows a node may be requested to handle.

Q30: For SCE, to receive credit for packet delivery, must packets be delivered in-order?

A30: No.

Q29: For SCE, will two transmit and two receive antennas be available in all matches?

A29: Yes.

Q28: For SCE, will the same *mandated_outcomes.json* file be supplied to all nodes within a network?

A28: Yes.

Q27: For SCE, will all **scored IP traffic flows have a corresponding mandated outcome?**

A27: Yes.

Q26: Per Q6, is it acceptable if a CIRN submission meets the specified total throughput, but fails to deliver all flows?

A26: Yes.

Q25: Will all scored IP-traffic flows use port numbers greater or equal to 5000?

A25: Yes.

Q24: For SCE, if the IP traffic flow persists across stage boundaries, can steady state period (SSP) times be met across stage boundaries?

A24: Yes.

Q23: When *scenario_center_frequency* in *environment.json* changes during a match, does this value supersede the value of *center_frequency* set in *colosseum_config.ini*?

A23: Yes.

Q22: For SCE, is the parameter *goal_set* in *mandated_outcomes.json* relevant to the score?

A22: This parameter is used only by the visualization engine and has no bearing on the score.

Q21: For SCE, will IP traffic flows include TCP traffic?

A21: No.

Q20: For SCE, will IP traffic flows include fragmented packets?

A20: No.

Q19: For SCE, how will file transfers be modeled?

A19: File transfers are modeled as a short burst of UDP packets whose aggregate payload is equal to the original file size.

Q18: Per the SC2 Phase 3 Scoring Procedures, what denotes the start of the earliest Measurement Period (MP) in a match?

A18: The earliest time that demarks the start of a Measurement Period is the official scenario start time provided by Colosseum in the match metadata.

Q17: Is a successfully delivered packet scored during the Measurement Period (MP) in which it originated, or the MP in which it was delivered?

A17: The originating MP.

Q16: For SCE, what is the minimum expected value for *max_latency_s* per the SC2 Phase 3 Scoring Procedures, Section 3.2 Individual Mandates?

A16: 100ms.

Q15: For SCE, must all packets in a flow with a mandated outcome which specifies only a *file_transfer_deadline_s* performance threshold (PT) be delivered to achieve the Individual Mandate (IM)?

A15: Yes.

Q14: Will DARPA release a script or other executable scoring “program” which computes a match’s score per the SC2 Phase 3 Scoring Procedures?

A14: No.

Q13: For SCE, what is the required number of nodes in a CIRN?

A13: 10.

Q12: Is the field *hold_time* in MandatedOutcomes.json equivalent to the Steady State Period (SSP) in the SC2 Phase 3 Scoring Procedures?

A12: Yes

Q11: Can we rely on the flow ids in the mandated outcome JSON file matching the destination port number of the packets for SCE?

A11: Yes.

Q10: Is a CIRN permitted to transmit messages over the collaboration network identifying itself as an incumbent or a different team?

A10: No.

Q9: For SCE, will any scored matches disable the collaboration network?

A9: No.

Q8: Are competitors wholly responsible for selecting all USRP settings, such as amplifier gain?

A8: Yes, each team’s CIRN software is responsible for configuring all USRP settings. For convenience, a list of recommended USRP settings has been provided by the USRP hardware manufacturer (National Instruments) and is available at the following location:

<https://sc2colosseum.freshdesk.com/support/solutions/articles/22000220403-optimizing-srn-usrp-performance>

Q7: Per Section 3.2.1 of the SC2 Rules Document, what non-performance related criteria must my CIRN achieve in order to qualify to participate in SCE?

A7: In order for a submitted CIRN design to qualify to participate in SCE, the CIRN must be compliant with the CIRN Interaction Language (CIL). Non-compliant submissions may be deemed ineligible. Compliance is determined using the CIL Validation Procedure available at <https://gitlab.com/darpa-sc2-phase3/CIL/blob/master/doc/CIL-Validation-Procedure.md>.

Q6: Per Section 3.2.1 of the SC2 Rules Document, what throughput performance must my CIRN achieve in order to qualify to participate in SCE?

A6: In order for a submitted CIRN design to qualify to participate in SCE, the 10-node CIRN must achieve the following specified aggregate data rates and latencies, using a 10MHz bandwidth, under the specified conditions:

SNR	Aggregate data rate achieved	Per packet latency achieved
5 dB	5 Mb/s	1.00 sec
10 dB	10 Mb/s	0.75 sec
15 dB	15 Mb/s	0.50 sec
20 dB	20 Mb/s	0.37 sec

This criteria will be tested in the following configuration:

- Each SNR condition will last 2 minutes
- The aggregate data rate must be consistently achieved over any consecutive 60 second period within the allotted 2 minutes per SNR
- The RF conditions will present the same approximate SNR between all nodes of the network
- UDP traffic given to each node will be sufficient to meet the required objective
- Every node in the network will be given the same offered load of data to transmit
- SNR given assumes a 10MHz bandwidth

Submissions which transmit outside the 10MHz of bandwidth (per the criteria outlined in the SCE Scoring Procedures Document) may be deemed ineligible.

Q5: During SCE, what is the maximum available frequency bandwidth a CIRN may transmit over without incurring a penalty for transmitting on a disallowed frequency?

A5: 40 MHz.

Q4: For SCE, what center frequency range will SRNs use?

A4: 900 MHz - 1100 MHz.

Q3: Will our CIRN be told how long a match will last?

A3: No. Information about match duration will not be provided to CIRNs.

Q2: May teams add or change members as the competition evolves through the phases?

A2: Yes, with the following restrictions:

- 1) A team may remove members at any point in the competition at the sole discretion of the team.
- 2) A team may add wholly new members at any point at the sole discretion of the team. Wholly new members are those that have not previously participated on any SC2 team.
- 3) A team may only add new team members who previously participated on another SC2 team under one of the following conditions:
 - The new team member participated as a member of another team in a previous phase of SC2, and has not participated in the current phase.
 - The new team member participated as a member of another team in a previous phase, and is changing teams during the allowed grace period, January 1 – March 31 of the next phase.
 - The new team member participated as a member of another team which has formally disbanded or otherwise withdrawn from the competition.
- 4) Teams may not add team members who currently or previously participated as part of the SC2 DARPA team.
- 5) All team membership changes must be registered with DARPA SC2 Team via email to sc2@darpa.mil.
- 6) The registered team lead is the only person authorized to make membership changes.

Q1: During the competition, can we (humans) provide control input to our nodes?

A1: No. CIRNs must be fully autonomous.