STRATEGIC
MONOGRAPH SERIES

OPERATION EPIC FURY:
A MULTI-DOMAIN FORESIGHT ASSESSMENT

BY DR. FREDERIC LEMIEUX AND DR. SHADI ABOUZEID



VOLUME 07/2026
RELEASE: JULY 2026

STRATEGY INTERNATIONAL

All content included in this publication is the property of
Strategy International (SI) Ltd.
Protected from the copyright laws of the Republic of Cyprus and international copyright laws.



Strategy International (SI) Ltd, Copyright ©2022-2027

Strategy International (SI) Ltd is a registered company HE423632 In the Republic of Cyprus

The logo of Strategy International is a registered trademark in the Republic of Cyprus. The trademark and works of strategy international (SI) Ltd and its logo are applied for class 31&45, in the Republic of Cyprus and EU IPO office.

Cyprus copyrights are governed by the Copyright Section of the Department of the Registrar of Companies and Official Receiver.

The governing laws are: N.63/77, N.18(I)/93, N.54(I)/99, N.12 (I)/2001, N.128 (I)/2002, N.128 (i) 2004 and N.123 (I) 2006.

Copyright protection vests automatically in the Republic of Cyprus.

First Published July 2022
Address: 24 Strovolos Str. 2042 Nicosia Cyprus

This publication is the exclusive property, including copyrights, of the author and Strategy International (SI) Ltd.

This publication in its complete form, or any parts of it, may not be reproduced, duplicated, copied, sold, resold, visited, or otherwise exploited for any commercial purpose. For all purposes, there is a need for a first written consent of the author and the legal representative of Strategy International (SI) Ltd.

SI trademark(s) may not be used in connection with any product or service and publication that has not been given any prior written consent first.

The written publication and opinions are considered of scientific and professional value and as such a result of the ongoing work of the author.

Any official statement made or position through all publications, interviews and related communication does not necessarily reflect the mission, objective and works and official opinions of Strategy International (SI) Ltd.

The document produced is endorsed, is its scientific value, and credits are provided to the author.

The editorial Team of
Strategy International (SI) Ltd.

Marios P. Efthymiopoulos (PhD) (CEO)
Theofaneia Ntounia (PhD cand)

Marketing Team: Rebel Online

Executive Summary

This report consolidates the analytical work of seven teams convened for *Operation Epic Fury*, a five-day scenario exercise (13–17 April 2026) conducted as part of the Global Security Futures course at Georgetown University's School of Continuing Studies. Working from a common fictional scenario, a joint United States–Israeli campaign against Iran and its consequences for the Gulf Cooperation Council (GCC) and the United Arab Emirates (UAE) in particular, each team examined a distinct domain along a shared methodological approach. A daily War News Brief advanced the scenario in real time, while structured analytic techniques, baseline and best- and worst-case scenarios, PESTEL and SWOT, red teaming, and the key-assumptions check, counterfactual "what-if" reasoning, and low-probability "wildcard" analysis, disciplined each judgment and tied it to a monitorable indicator. The scenario and its data are fictional; the methods applied, the structural dynamics they expose, and the regional vulnerabilities they illuminate are not.

The seven assessments converge on a single, sobering proposition: the war has not so much disrupted the Gulf security order as revealed that the order had already entered transition. Eight key judgments anchor the finding. The conflict has settled into a managed stalemate that neither belligerent can break at an acceptable cost. Leadership decapitation created a visibility gap for the coalition rather than a capability gap for Iran, whose decentralized command was designed to survive it. GCC cohesion is eroding competitively rather than collectively, with each member hedging individually, and Qatar's swing behavior is the single load-bearing variable. The region's physical vulnerabilities, the Strait of Hormuz, the Fujairah corridor, and coastal desalination, form a single interlocking "chokepoint stack" in which stress on one node transmits to the others. UAE strategic neutrality emerged as the structural element keeping the conflict limited, and its collapse is the highest-leverage downside trigger. The information environment has no prior baseline to revert to; economic damage is compounding, with recovery determined by preparation made during the conflict rather than after it; and an unconventional-weapons threshold event, though low-probability, remains the dominant tail risk.

Written from a vantage point six weeks after the 17 April analytical cutoff, the report finds that the teams' framework has aged well: the stalemate has held, the fragmentation

they warned of has accelerated, and the UAE's subsequent withdrawal from OPEC rendered the competitive-fragmentation thesis literal. The most actionable learning is that the stalemate is a pressure-accumulation process, not a resting point, and that the only variable fully within the GCC's control is the speed of its preparation. The present interval of relative quiet is therefore a window to be used, not a reprieve to be waited out. Each state's capacity, pre-positioned now, preserves agency, whereas waiting entails consequences. Beneath the seven domains lies one largely unnamed choice, whether the Gulf builds its emerging security order together or apart.

Finally, the exercise demonstrates that the enduring value of structured foresight lies less in the precision of any single probability estimate than in the discipline it imposes. Surfacing the assumptions beneath each judgment, tracing every alternative future to an observable signpost, and converting analysis into a monitoring program that decision-makers can actually run. The Consolidated Signpost Watch-List that unifies these assessments is the practical expression of that discipline, and the report's most durable contribution to those who must act amid uncertainty.

Acknowledgments

We wish to thank Strategy International (SI) Ltd, for the opportunity to share our thoughts and insights on the critical issue of the Iran Vs USA / Israel war and its implications across the GCC in general and UAE in particular.

We would also like to personally thank Strategy International CEO Dr. Marios P. Efthymiopoulos.

Table of Contents

Executive Summary iv
Acknowledgments vi
Table of Contents 7
Introductory Remarks by the Editors 10
Introduction 10
Methodology 12
The Daily Brief and the Evolving Scenario 12
Analytical Foundations: The Intelligence Cycle and the Fog of War 12
Baseline and Alternative Scenarios 13
Environmental Scanning: PESTEL and SWOT 13
Cognitive Bias and Red Teaming 14
Counterfactual "What-If" Analysis and Back-Casting 15
Wildcard Scenarios and the Fait Accompli 15
Synthesis, Publication, and Limitations 16
PART I: Strategic Assessment 17
Key Judgments 18
The Cross-Cutting Picture 21
The Decapitation Paradox 21
The Chokepoint Stack 21
Competitive Fragmentation 22
The Information War Without an Off-Ramp 22
The Load-Bearing Assumptions 23
Consolidated Signpost Watch-List 24
PART II: Domain Assessment 27
Iranian Intent and Command Coherence Under Decapitation 29
Situation and Baseline 29
Best and Worst Case 29
The Critical What-If 30
The Wildcard 30
Decision Implications 31
The Strait of Hormuz Under Dual Interdiction 32
Situation and Baseline 32
Best and Worst Case 32

The Critical What-If 33
The Wildcard 33
Decision Implications 33
The Energy–Water Nexus and Its Human Operating Layer 35
Situation and Baseline 35
Best and Worst Case 35
The Critical What-If 35
The Wildcard 36
Decision Implications 36
Hydrocarbons, Tourism, and the Speed of Recovery 37
Situation and Baseline 37
Best and Worst Case 37
The Critical What-If 37
The Wildcard 38
Decision Implications 38
Narrative Integrity and the Architecture of Belief 39
Situation and Baseline 39
Best and Worst Case 39
The Critical What-If 39
The Wildcard 40
Decision Implications 40
The Sustainability of the Hedging Posture 41
Situation and Baseline 41
Best and Worst Case 41
The Critical What-If 41
The Wildcard 42
Decision Implications 42
Cohesion, the Qatar Variable, and the Post-Conflict Order 43
Situation and Baseline 43
Best and Worst Case 43
The Critical What-If 43
The Wildcard 44
Decision Implications 44
PART III: The Outlook 45
The Stalemate Is a Process, Not a Resting Point 46
The Variables Most Likely to Break It 46

What the Window Is For 47
The Choice the Gulf Is Actually Making 47
A Closing Note on the Method 48

Introduction

As of 29 May 2026, three months after the opening strikes of the joint United States–Israeli campaign against Iran, Washington and Tehran have yet to conclude an agreement. The Strait of Hormuz, the maritime chokepoint through which roughly a quarter of the world's seaborne crude oil and a fifth of its liquefied natural gas ordinarily pass, continues to function as a contested and politically filtered passway. It is constrained simultaneously by Iranian control of its northern approaches and by a United States naval blockade of Iranian ports to the south. The economic consequences remain acute and global in scope. Brent crude, having traded near $72 per barrel before the conflict, has held above $110 for weeks; war-risk insurance for Gulf transits is either prohibitively priced or wholly unavailable; and the disruption has propagated well beyond petroleum. Because Gulf natural gas is a principal feedstock for the world's nitrogen-based fertilizers, constrained gas and LNG flows have tightened global fertilizer supply. The latter directly drove upward food prices, in addition to the wider commodities and manufactured goods that transit the Strait, now absorbing higher freight costs and longer routes. The conflict that many observers expected to resolve within weeks has instead consolidated into the very condition that the contributors to this volume identified, a month before its persistence was confirmed, as the single most probable trajectory: an unresolved, pressure-accumulating stalemate.

This report synthesizes the analytical work of seven teams convened for Operation Epic Fury as part of the Global Security Futures course within a program at Georgetown University's School of Continuing Studies. Over five days (13–17 April 2026), each team conducted a structured scenario exercise covering the impact of the war on the Gulf Cooperation Council (GCC) in general and the United Arab Emirates (UAE) in particular. It bears emphasis at the outset that the scenario, the events it describes, and the data underlying it are fictional and were constructed solely for educational purposes. The analytical methods employed, the structural dynamics they expose, and the regional vulnerabilities they illuminate are not.

The seven teams examined a common scenario from distinct vantage points but along a shared methodological spine. Each developed a baseline assessment of the most probable near-term future, expressed where appropriate in the calibrated probabilistic language of established intelligence-community analytic standards, and bounded that baseline with explicit best- and worst-case scenarios. Each then subjected its judgments to structured challenge through red teaming, PESTEL and SWOT analysis, the construction of "what-if" trigger pathways, and the deliberate examination of low-probability, high-impact "wildcard" events, to surface cognitive biases, test assumptions, and widen the aperture of strategic foresight. The common architecture of the assessments that follow is therefore a product of shared method rather than coincidence, and it permits a degree of cross-team comparison that single-domain analysis would not.

The domains span the principal dimensions of the crisis. The strategic-intelligence assessment interrogates Iranian intent and the coherence of Iranian command and control under conditions of leadership decapitation. The critical-infrastructure analysis examines the energy–water nexus and the predominantly expatriate workforce on which its operation depends. The maritime assessment addresses the governance of access to the Strait of Hormuz as great-power competition extends into Gulf waters. The economic and supply-chain analysis treats the hydrocarbon, tourism, and hospitality sectors. The cognitive and information security study assesses the contest over narrative integrity and decision-making within Gulf societies and institutions. The assessment of the UAE examines the sustainability of the Emirates' strategic hedging posture when its foundational infrastructure nodes are affected. Any major disruption to either the desalination plants or the Fujairah corridor might push the UAE leadership into a binary choice of alignment that neutrality is designed to avoid. Finally, the regional-stability assessment evaluates the durability of Gulf political cohesion and the contours of a prospective post-conflict order.

Read in aggregate, the seven assessments converge upon a single and uncomfortable proposition. The war has not so much disrupted the Gulf security order as revealed. The order was already in transition, and the member states of the GCC, confronted by a common adversary, have for the most part responded competitively rather than collectively. The analysis that follows seeks to render that transition legible while the opportunity to influence its direction remains open.

Methodology

The seven assessments collected in this volume were produced under a shared methodological protocol across the five days of the exercise (13–17 April 2026). Each day combined three elements in a fixed sequence: a morning intelligence brief updating the scenario, a lecture introducing a specific analytic method, and an extended afternoon session in which the teams applied that method to their domains. The analytic products accumulated rather than replaced one another. Teams first established a baseline and its best- and worst-case boundaries; they then subjected that baseline to structured adversarial review; next, they constructed counterfactual "what-if" analyses; and finally, they developed low-probability, high-impact "wildcard" scenarios. By the close of the week, each team held a layered assessment in which every alternative future was tied back to the baseline it tested and forward to an indicator that could be monitored in the present. This section describes the protocol and the structured analytic techniques on which it rested.

The Daily Brief and the Evolving Scenario

A defining feature of the exercise was that the scenario did not stand still. Each morning opened with a War News Brief: an intelligence update introducing new developments, incidents, and reporting that had "occurred" since the previous session. These daily injections serve the same function as real-world reporting for a standing analytic cell. They introduced fresh and often contradictory information, advanced or reversed trends the teams had been tracking, and forced every team to decide whether each new development confirmed its prior assessment, warranted a revision, or could be set aside as noise. Because the briefs arrived in sequence over a compressed timeline, they reproduced, in miniature, the central problem of conflict analysis: that assessments must be updated continuously against a moving target, under time pressure, and without the luxury of waiting for ambiguity to resolve. The daily brief was therefore not merely a narrative device but the mechanism that made the exercise dynamic. It ensured that each day's analytic technique was applied to a situation that had materially changed since the previous one. In addition, it created the conditions under which anchoring, confirmation bias, and sunk-cost reasoning, the failure modes the red-teaming sessions were designed to surface, could actually take hold and be observed.

Analytical Foundations: The Intelligence Cycle and the Fog of War

The exercise opened by establishing the analytic environment rather than an analytic method. The first day situated each team within the five-phase intelligence cycle: planning and direction, collection, processing, analysis, and dissemination. This step highlighted the ways each phase fails under wartime tempo. Consider the case of requirements that are obsolete by the time they are answered, degraded or denied collection, processing bottlenecks in which raw reporting arrives faster than it can be exploited, premature closure at the analytic stage, and last-mile failures in which finished assessments never reach decision-makers in time. Over this was laid the "Clausewitzian" distinction between fog, uncertainty about the operating environment, and friction, the accumulation of small frustrations that resist the execution of

any plan. These concepts were not abstract: the scenario confronted each team with concrete instances of both, from unverifiable air-defense intercept claims to genuine ambiguity over whether Iranian strikes on Gulf civilian infrastructure were centrally directed, crisis-driven, or the product of autonomous provincial action.

Two features of the scenario shaped every subsequent assessment. The first was leadership decapitation: the coalition's targeting of the Supreme Leader and the installation of a contested successor. The second, and analytically decisive, was Iran's decentralized resilience architecture, a model in which each of thirty-one provinces holds pre-positioned weapons, independent intelligence networks, self-sufficient logistics, and, crucially, pre-delegated authority allowing provincial commanders to continue operating without orders from Tehran. The governing insight, introduced on the first day and borne out across the team reports, was that against such an architecture, decapitation produces a visibility gap for the attacker rather than a capability gap for the adversary. The coalition loses the single interlocutor capable of authorizing de-escalation precisely as autonomous operations continue. This proposition recurs, under different names, in the strategic-intelligence, maritime, infrastructure, and UAE assessments, and it explains why several teams treated the question of whether Iranian behavior was centrally directed as more consequential than the content of any Iranian decision.

Baseline and Alternative Scenarios

With the environment established, the teams built their baseline assessments: a judgment of the single most probable near-term trajectory in each domain, accompanied by an explicit list of the core assumptions on which it depended. Each baseline was bounded by a best case and a worst case, and each boundary scenario was specified not as a mood but as a set of enabling conditions, for the best case, or triggers, for the worst. Confidence was expressed in calibrated probabilistic language drawn from intelligence-community analytic standards, the ICD-203 likelihood bands in some teams' work, explicit percentage ranges in others, so that judgments could be compared and, later, revised. The requirement to surface assumptions explicitly was deliberate: an assumption that is written down can be challenged, reversed, and monitored, whereas an assumption that remains implicit functions, in the language of the exercise, as an invisible anchor.

Environmental Scanning: PESTEL and SWOT

To discipline the scan of the external landscape, the teams applied two complementary frameworks. PESTEL (political, economic, social, technological, environmental, and legal) provided a structured checklist of macro-environmental forces, ensuring that no dimension of the operating environment was overlooked. In addition, cross-pillar interactions, such as an economic shock with social and legal consequences, were made visible rather than missed. SWOT then mapped each actor's internal strengths and weaknesses against the external opportunities and threats identified in the PESTEL scan, and the associated TOWS matrix

converted those static inventories into candidate strategies, offensive, investment, defensive, and survival. Factors were scored on two dimensions, typically impact and likelihood, to force prioritization rather than undifferentiated listing. Applied to the Gulf Cooperation Council, this produced an honest position assessment, substantial sovereign-financial capacity and strategic geography on one side, direct exposure of ports, pipelines, and desalination plants on the other, that several teams carried forward into their baselines and decision recommendations.

Cognitive Bias and Red Teaming

The methodological core of the exercise was the deliberate, structured challenge of the teams' own judgments. The instruction grounded this in dual-process psychology: the fast, automatic reasoning that generates intuitive judgments, together with the systematic biases that accompany it, and the slow, effortful reasoning that can correct them but rarely engages on its own. Structured analytic techniques were presented as forcing functions for the latter. Seven biases were treated as the principal hazards of conflict analysis: anchoring, confirmation bias, the availability heuristic, mirror imaging, groupthink, sunk-cost reasoning, and overconfidence, with mirror imaging singled out as the most dangerous because it is the least visible. The analyst who projects Western rational-actor logic onto Iran, assuming, for instance, that economic pressure produces a move toward negotiation or that decapitation paralyzes command, makes no active error and so has nothing to detect. Against an adversary whose decentralized architecture was explicitly designed to invert those assumptions, mirror imaging was not a generic caution but a specific and recurring threat, and the exercise's bias audit flagged it by name for the maritime and UAE teams, among others.

Red teaming was defined narrowly and against its common misreadings: not criticism, not contrarianism, and not unstructured brainstorming, but organized skepticism conducted under explicit permission, focused on assumptions rather than personalities, and documented even when its challenges were ultimately rejected. Six techniques were available: the devil's advocate, analysis of competing hypotheses, the premortem, the key assumptions check, outside-in analysis, and Team A / Team B analysis, each matched to a particular failure mode. In practice, the teams worked primarily from a key assumption check. Each team enumerated at least eight explicit and implicit assumptions in its baseline, classified them by criticality and falsifiability, and selected the three most load-bearing. For each assumption, the teams specified the evidence that would break it, the alternative hypothesis that would then emerge, and the probability estimate that would follow. Teams reviewed one another's work rather than only their own, on the principle that adversarial review is most effective when conducted from outside the originating team. The effect was substantial and is visible in the "effect of red teaming" discussion in each report. Baselines were generally made more pessimistic and more provisional. Several high-confidence judgments were downgraded, and a number of assumptions that had been carrying decisive weight, that Iranian behavior was centrally directed, that opacity in a sovereign wealth fund implied fragility, that a particular

diplomatic framework was a settled fact rather than single-source reporting, were reclassified once stated and tested.

Counterfactual "What-If" Analysis and Back-Casting

The third method extended and stress-tested the baseline through counterfactual, or "what-if," analysis, framed by escalation dynamics and deterrence theory. A counterfactual was defined as structured causal reasoning from a single altered premise. One variable was reversed at a time, the premise had to be genuinely plausible rather than imaginative, and it was built in stages. It precisely defined the altered state before any consequences unfolded, traced the first-order effects within forty-eight to seventy-two hours, followed those effects across domains as military consequences became economic and economic consequences became diplomatic, and concluded with a forward inference. That final step was treated as non-negotiable: a what-if that does not end in something an analyst can monitor in the present was deemed history rather than intelligence. The companion technique, back-casting, reversed the direction of reasoning, beginning from a concrete scenario endpoint and working backward through what would have to be true one week before, one month before, and at the earliest observable signal. This allows the conversion of each alternative future into a prioritized watchlist of signposts, which is assigned to a team responsible for monitoring it. The analytic payoff was leverage: by reversing the United Arab Emirates' decision to recall its ambassador, for example, the exercise demonstrated that Emirati neutrality carried more causal weight than almost any other political variable, which in turn identified monitoring that neutrality as a first-order priority.

Wildcard Scenarios and the Fait Accompli

The fourth method addressed the low-probability, high-impact events that, by construction, fall outside both the baseline and the standard worst case. The instruction distinguished three kinds:

1. The black swan, genuinely outside the range of expectation, whose failure mode is imagination.
2. The known wildcard, acknowledged on the risk register but under-monitored, whose failure mode is attention rather than imagination.
3. The accelerant, which does not change a trajectory's direction but compresses its timeline, forcing decisions before assessments can be updated.

Wildcards were defined as living in the shared blind spots of all the teams, the developments that every team's analysis quietly assumed would not happen and were to be selected for plausibility rather than probability, on the test of whether the event could occur and whether its approach would be recognized. Each wildcard was then developed by tracing three distinct fallout pathways and, for each, identifying specific, monitorable signposts assigned to a

responsible team. A complementary “fait accompli” exercise inverted the analytic stance. Teams were presented with a transformative event, shown below, as an accomplished fact:

- an internal IRGC fracture
- a sovereign-wealth-fund liquidation triggering a global liquidity crisis
- a Chinese carrier deployment into the Gulf of Oman
- a viral deepfake of the Iranian leadership ordering chemical-weapons use

Teams were tasked with mapping its consequences rather than assessing its likelihood, on the premise that for such events, the analytically useful question is not whether they will occur but whether the analytic architecture would absorb them if they did.

Synthesis, Publication, and Limitations

On the final day, the teams prepared their assessments for publication, which accounts for the common architecture of the chapters that follow:

- a baseline, best- and worst-case boundaries
- a what-if
- a wildcard
- a concluding decision implication
- where possible a monitorable indicator.

Several limitations of the method should be stated plainly. The scenario and its data were fictional and compressed into five days, and the near-term horizons the teams adopted varied by domain, so the probability estimates are internally calibrated rather than externally validated. The analysis relied on open-source and scenario-provided information without access to classified collection or live telemetry. And the exercise was conducted once, without the independent replication that the Team A / Team B technique models. These constraints bound the claims that follow: the value of the assessments lies less in the precision of any single probability estimate than in the structured identification of assumptions, leverage points, and early-warning indicators that decision-makers can monitor as the actual situation evolves.

PART I: Strategic Assessment

Information cutoff: This assessment reflects information available to the analytical teams through the close of the exercise on 17 April 2026, and the probability estimates and judgments below are anchored to that date. This report is published from a later vantage point (29 May 2026); where a development after the cutoff bears materially on a judgment, it is identified as a dated editorial update and kept separate from the underlying assessment. All statements not so marked should be read as of 17 April 2026.

Key Judgments

KJ-1. The war has settled into a managed stalemate that neither belligerent can break at an acceptable cost. ***(High confidence.)*** As of the cutoff, the conflict had stabilized into a mutual-denial equilibrium. The United States imposed a naval blockade of Iranian ports to the south. Iranian control of the Strait of Hormuz's northern approaches to the north. Neither side can force a resolution without unacceptable escalation. The strategic-intelligence, maritime, and economic assessments independently judged this condition the most probable near-term trajectory rather than a transitional phase, principally because the structural features sustaining it, decentralized Iranian command and a coalition unwilling to mount a ground campaign, were themselves stable.

KJ-2. Leadership decapitation created a visibility gap for the coalition, not a capability gap for Iran. ***(High confidence.)*** The targeting of the Supreme Leader and the contested succession degraded the coalition's insight into Iranian decision-making far more than it degraded Iran's ability to act. Iran's decentralized resilience architecture, provincial commands holding pre-positioned weapons, independent collection, and pre-delegated authority, was designed precisely to survive decapitation. The recurring analytical hazard, flagged across at least three teams, is the mirror-imaging assumption that removing the head halts the body; the more dangerous reality is that autonomous operations continue while the one actor capable of authorizing de-escalation has been removed.

KJ-3. GCC cohesion is eroding competitively, not collectively. ***(Moderate-to-high confidence.)*** Confronted by a common threat, the member states are responding as individual hedgers rather than as a bloc, each optimizing its own exposure, alignment, and economic continuity. Gulf unity is not collapsing in a single rupture; it is being quietly outcompeted from within. Qatar's willingness to co-sign collective positions emerged as the single load-bearing variable; its swing behavior, more than any formal mechanism, determines whether the GCC acts as one.

KJ-4. The Gulf's physical vulnerabilities form a single interlocking "chokepoint stack," not a set of separate risks. ***(High confidence.)*** The Strait of Hormuz, the Fujairah/ADCOP bypass corridor, the Bab el-Mandeb approach, and coastal desalination capacity are analytically inseparable. Stress on any one of them transmits to the others, and the bypass routes intended to mitigate Hormuz dependence were themselves shown to lie within the battlespace. The maritime, economic, infrastructure, and UAE assessments converge on the

judgment that resilience planning premised on any single alternative route is planning on a node that can also be struck.

KJ-5. UAE strategic neutrality is the structural element keeping the conflict limited; its collapse is the highest-leverage downside trigger. ***(Moderate-to-high confidence.)*** The Emirates' tolerated-but-unendorsed hedging posture preserves the only viable diplomatic off-ramp and caps the conflict's geographic spread. That posture rests on a narrow base composed of the integrity of desalination and the Fujairah corridor. A forced binary alignment, whether through a strike on those nodes or through coalition pressure, would remove a stabilizing variable carrying more weight than almost any other political factor in the scenario.

> **Editorial update (as of 29 May 2026):** *A development after the cutoff bears directly on this judgment. On 28 April 2026, eleven days after the exercise closed, the UAE announced its withdrawal from OPEC and OPEC+, effective 1 May, becoming the largest producer ever to leave the organization after nearly six decades of membership. officials framed it as a sovereign long-term energy policy decision and stated they had not consulted Saudi Arabia or other members beforehand. The move cuts in two directions. It does not contradict the core finding; the UAE continues to pursue an independent, interest-first course rather than bloc alignment. Several analysts read the exit as a visible marker of a Riyadh–Abu Dhabi rupture, which is the competitive-fragmentation dynamic of KJ-3 made concrete. But it complicates the reading of UAE behavior as conflict-limiting "neutrality": an unconsulted, unilateral break with the principal Gulf production framework is a more assertive and divisive act than the passive hedging assessed through 17 April. It also suggests an erosion in the intra-GCC cohesion and less quietly, than the cutoff-date analysis implied. The monitoring implication is unchanged but sharpened; UAE alignment signals remain the highest-priority political indicator, while the direction of travel since the cutoff has been toward fragmentation.*

KJ-6. The information environment has no analog baseline to revert to. ***(Moderate confidence.)*** Narrative contestation across Gulf publics and institutions is now a permanent operating condition rather than a discrete campaign with a beginning and end. Because the underlying information and decision systems are themselves the contested terrain, and because attempts at information control were assessed as frequently self-defeating, there is no prior, lower-tempo state to which the environment can simply return once kinetic operations subside.

KJ-7. Economic damage is compounding, and the decisive variable is the speed of recovery preparation, not the depth of the shock. ***(Moderate confidence.)*** Sustained elevated oil prices, suspended war-risk insurance, and disrupted hydrocarbon, tourism, and hospitality flows were assessed as recoverable in principle, but the pace of recovery depends on preparations made during the conflict rather than after it. The states that emerge strongest will be those that pre-positioned redundancy, storage, and reconstruction capacity while the war was still ongoing.

KJ-8. A nuclear- or chemical-threshold event is low-probability but is the dominant tail risk. ***(Low confidence / high impact.)*** Across the wildcard analyses, the events capable of restructuring every team's assessment overnight clustered around the unconventional-weapons threshold, proliferation pressure, a demonstrative detonation, or a chemical-use allegation, whether genuine or fabricated. None was assessed as likely, but each would be transformative, and each is therefore retained as a monitored tail risk rather than excluded for low probability.

The Cross-Cutting Picture

The eight judgments above are not eight separate findings. Read together, the seven assessments describe a single system under stress, and the same structural features surface in domain after domain under different names. This section draws those convergences out explicitly. It organizes the material into five cross-cutting themes, the decapitation paradox, the chokepoint stack, competitive fragmentation, the information war without an off-ramp, and the load-bearing assumptions on which the entire picture rests. Each theme is the product of several teams' work rather than any one, and together they form the analytical core of this report; the domain chapters in Part II supply the evidence behind them.

The Decapitation Paradox

The most consequential finding of the exercise is also its most counterintuitive: the coalition's greatest tactical success, the removal of Iran's supreme leadership, produced the condition that most complicates its strategic position. The reasoning runs through three steps that recur across the strategic-intelligence, infrastructure, and maritime assessments. First, Iran's command architecture was decentralized by design: each of thirty-one provinces holds pre-positioned weapons, independent intelligence networks, self-sufficient logistics, and pre-delegated authority to act without instruction from Tehran. Second, against such an architecture, decapitation removes the apex of a hierarchy that the lower levels do not actually depend upon to function, so Iranian operations continue, dispersed and uncoordinated, rather than halting. Third, and decisively, the actor removed is the only one with the authority to order de-escalation; eliminating him therefore subtracts the off-ramp while leaving the vehicle in motion.

The paradox is what this does to the coalition's knowledge. Decapitation opened a visibility gap rather than a capability gap: the coalition lost its clearest read on Iranian intent precisely as autonomous provincial action made that intent harder to infer from behavior. A strike on Gulf civilian infrastructure can no longer be read confidently as a centrally directed coercive signal, because it may equally be an uncoordinated provincial initiative, and the two demand opposite responses. This ambiguity is not noise to be resolved; it is the central analytical condition of the conflict, and it is why several teams treated the question of whether Iranian behavior was centrally directed as more important than the content of any single Iranian act. The corollary is a standing warning against mirror imaging: the intuition that pressure produces recalculation, or that decapitation produces paralysis, is precisely the Western rational-actor assumption the Iranian architecture was built to defeat.

The Chokepoint Stack

The exercise's second convergence concerns geography. Analysts across the maritime, economic, infrastructure, and UAE assessments stopped treating the Gulf's physical vulnerabilities as a list of discrete risks and began treating them as a single interlocking system. A "chokepoint stack" in which the Strait of Hormuz, the Fujairah and ADCOP bypass routes, the Bab el-Mandeb approach, and coastal desalination capacity are coupled tightly enough

that stress on any one propagates to the rest. The Strait is the obvious node, carrying roughly a quarter of seaborne crude and a fifth of LNG; but the bypass infrastructure built specifically to reduce Hormuz dependence, Saudi Arabia's East–West pipeline, the UAE's Fujairah corridor, was itself shown to lie inside the battlespace, which collapses the distinction between primary route and redundancy on which Gulf resilience planning has long relied.

Desalination is the element that makes the stack existential rather than merely economic. Gulf states depend on desalinated seawater for the majority of their potable supply. The plants are co-located with the coastal power and export infrastructure already under threat, and they are operated by a workforce that is overwhelmingly expatriate, a human layer whose departure under sustained attack would disable the plants as surely as a physical strike. The infrastructure and UAE assessments converge on the judgment that desalination is the true keystone: its loss does not degrade an economy; it empties a coastline. The strategic implication, picked up in KJ-4 and KJ-7, is that redundancy premised on any single alternative node is fragile, and that the meaningful resilience question is not which route survives but how quickly multiple simultaneous failures can be absorbed and repaired.

Competitive Fragmentation

The third theme is the one that gives the report its title proposition: the Gulf Cooperation Council is not failing to cohere so much as cohering competitively, each member pursuing an individually rational hedge that is collectively corrosive. The strategic logic is straightforward and sound from each capital's vantage point. In a conflict that threatens every member's export, alignment, and economic continuity, the dominant strategy is to minimize one's own exposure rather than to bind oneself to a common position that may carry someone else's risk. The result is not a dramatic rupture but a quiet divergence, separate channels to Tehran, separate calculations on alignment, separate bids to position as the indispensable post-conflict hub.

Two actors carry the theme. The regional stability and UAE assessments identify Qatar's willingness to co-sign collective positions as the single load-bearing variable for GCC unity: when Qatar holds, the bloc can present a common face; when it opens an independent channel, the common face dissolves. The Emirates is the other pivot; its hedging posture both stabilizes the conflict (KJ-5) and, paradoxically, exemplifies the fragmentation, since a successful individual hedge is itself a defection from collective discipline. The tension between those two readings of UAE behavior is not a contradiction to be resolved but the precise shape of the dilemma: the same posture that limits the war also erodes the alliance. The post-cutoff OPEC withdrawal, discussed under KJ-5, is this dynamic made unmistakable.

The Information War Without an Off-Ramp

The fourth convergence comes principally from the cognitive and information-security assessment but is reinforced wherever the other teams touch perception, attribution, or market psychology. Its central finding is structural rather than tactical. The contest over

narrative is no longer an episode layered atop the “real” conflict but a permanent condition of the environment, because the information and decision-making systems through which Gulf societies operate are themselves the contested terrain. There is no lower-tempo analog state to which the environment reverts when the missiles stop, because the vulnerability is the digital and cognitive substrate itself, which does not switch off.

Two observations sharpen the theme. First, attempts at information control were repeatedly assessed as self-defeating, censorship and heavy-handed narrative management tended to corrode the credibility of the controlling authority faster than they suppressed the hostile narrative, so the instinctive state response made the underlying problem worse. Second, the same artificial-intelligence capabilities that now pervade the information space make fabrication cheap and verification slow: a convincing forgery can circulate globally within hours, while authentication takes days, leaving decision-makers with windows in which they must act on information they cannot yet confirm. The combination, a permanently contested substrate, self-defeating control reflexes, and a widening gap between fabrication and verification, is why KJ-6 treats narrative integrity as a standing operating condition rather than a problem with a solution.

The Load-Bearing Assumptions

The final theme is methodological in origin but strategic in consequence. Because every team surfaced and tested the assumptions beneath its baseline, it is possible to identify the handful of propositions on which the entire assessment rests, the points where, if the assumption broke, multiple judgments would move at once. Four stand out. The first is that Iranian operations remain decentralized and pre-delegated rather than re-centralizing under the successor leadership; almost every judgment about escalation, attribution, and de-escalation depends on it (KJ-1, KJ-2). The second is that the Strait of Hormuz remains selectively passable rather than fully closed or fully open; the economic and maritime baselines are built on this middle condition, and either extreme would invalidate them (KJ-1, KJ-4). The third is that UAE neutrality holds; its failure is the highest-leverage single trigger in the scenario (KJ-3, KJ-5). The fourth is that no unconventional-weapons threshold is crossed; this assumption is rarely stated because it is rarely doubted, which is precisely what makes it dangerous (KJ-8).

The value of naming these four is that they convert a static assessment into a monitoring program. Each is an assumption whose reversal has been traced to specific, observable consequences, and each therefore points to a set of indicators that would give early warning that the picture is changing. That is the function of the consolidated watchlist in the section that follows. It gathers the indicators attached to these load-bearing assumptions, the wildcard and what-if analyses, and all seven domains into a single instrument that tells a decision-maker not merely what the teams concluded but what to watch for as evidence that the conclusions no longer hold.

Consolidated Signpost Watch-List

The seven assessments did not end in conclusions; they ended in indicators. Each team was required to convert its judgments into observable, open-source, monitorable signals, signposts, and to specify, for each, what its appearance would mean and how it should change the assessment. This section consolidates those signals into a single instrument. It is the operational core of the report: the place where seven separate analyses converge into a single monitoring program.

The table that follows is organized into three warning tiers. Critical-tier indicators are those whose appearance would, on its own, force an immediate revision of one or more key judgments and, in several cases, allow only hours to respond; these warrant continuous, named-owner monitoring with a pre-delegated response rather than periodic review. High-tier indicators would materially shift a probability estimate and warrant daily monitoring against the morning brief. Moderate-tier indicators are slower-moving structural signals to be reviewed on a regular cadence. For each signal, the table records the domain that owns it, what its appearance would signify, and the specific revision it should trigger. A signpost is not a prediction: its appearance raises the probability that a scenario is materializing, it does not confirm it, and a cluster of signposts across tiers should be treated as an analytical alert in its own right.

> *All indicators below are drawn from the seven domain assessments, finalized on 17 April 2026, and are designated as monitorable from that date forward. Ownership is assigned to the domain team best positioned to observe the signal; in practice, several indicators warrant cross-team monitoring, and those cross-references are noted where specified.*

Owner	Indicator (open-source observable)	Signals	If observed, revise
CRITICAL: continuous monitoring; appearance forces immediate revision (hours–days of warning)			
T1 / T2	Any autonomous provincial IRGC strike that contradicts Tehran's stated position, a kinetic action during a declared central de-escalation or ceasefire posture.	Confirms	The most load-bearing assumption in the assessment, that Iranian operations remain directable from the center, is falsified. Treat decentralized Mosaic-Defense behavior as the baseline (KJ-2); escalation, attribution, and de-escalation judgments revised accordingly.
T6 / T2	Iranian state media naming UAE water/desalination or pipeline infrastructure as a legitimate target.	Confirms	The UAE infrastructure what-if ceases to be a what-if and becomes the baseline. Activate desalination and Fujairah-corridor contingency posture; the conflict-limiting reading of UAE neutrality (KJ-5) is in immediate jeopardy.
T6 / T2	ADNOC or Fujairah Port Authority force majeure declaration.	Confirms	The decisive, observable point of no return in the UAE export-cascade pathway, there is no intervention point between this declaration and cascade completion. Treat the worst-case economic trajectory as active (KJ-4, KJ-7).

Owner	Indicator (open-source observable)	Signals	If observed, revise
T5	Coordinated suppression of analog/satellite emergency communications plus interference against commercial satellite terminals in the Gulf, preceding any kinetic event.	Confirms (begun)	This is not warning that the EMP/cable wildcard is approaching, it is confirmation it has begun. Initiate the pre-delegated, named-owner response within minutes; decision window is hours, not days.
T4	Resumed vehicular activity at Isfahan; a credible Iranian weapons-grade-enrichment claim; satellite detection of a test-preparation site; or any IAEA-affiliated government confirming material diversion.	Confirms	The unconventional-weapons tail risk (KJ-8) is materializing. Trigger the pre-delegated nuclear-signaling protocol without a new decision cycle; breakout windows are measured in days.
T7	Qatar delays co-signature on any GCC security statement by more than 48 hours.	Confirms	Revise worst-case GCC-fragmentation probability immediately from 25% to 45% (KJ-3); update Pillar I viability, Pillar II energy-escalation risk, and Pillar IV coordination resilience simultaneously.
HIGH: daily monitoring against the morning brief; materially shifts a probability estimate			
T3	Confirmed kinetic engagement against a flagged Indian, Chinese, Turkish, or Pakistani vessel inside Iran-designated Hormuz lanes.	Breaks	The assumption that Iran retains C2 coherence sufficient to enforce a discriminatory passage regime is broken; the selective-passage baseline (KJ-1, KJ-4) no longer holds.
T3 / T6	Lloyd's reinstating standard war-risk coverage for Hormuz transit, or more than 20 vessels transiting successfully in a 48-hour window.	Breaks (upside)	The "functionally closed Strait" assumption is falsified on the upside; economic and UAE baselines premised on closure should be re-rated toward recovery (KJ-1, KJ-7).
T3	Unscheduled PLAN Type 903A replenishment transit toward Bandar Abbas, or Chinese-flagged roll-on/roll-off arrivals at Iranian ports.	72-hr warning	Early indicator of Chinese material involvement reframing Beijing as co-belligerent in U.S. perception, the carrier-strike what-if pathway. GCC neutrality cannot survive that reframing.
T3	PLAN surge departure of Type 052D/054A/903A outside the Aden rotation; MFA "sea-lane security" language tied to Hormuz; AIS grouping of Chinese LNG carriers in Gulf of Oman anchorages; Chinese port-access inquiries at Khalifa, Fujairah, Duqm, or Jebel Ali.	Weeks warning	Precursors to a Chinese Hormuz escort declaration, the structural end of single-patron Gulf maritime order. Begin renegotiation planning for a two-patron access environment.
T6	OFAC investigations into UAE banking flows; congressional hearings naming UAE banks; Chinese state banks visibly expanding UAE operations; sovereign CDS spreads widening.	Weeks warning	Precursors to a U.S. Treasury secondary-sanctions ultimatum on UAE banks, the one pressure the hedging posture cannot survive (KJ-5). Compress the decision window via pre-filed compliance audit and parallel settlement track.
T1	Hezbollah escalation beyond the ceasefire expiry, Houthi Bab el-Mandeb activation, or IRGCN "parting-shot"	Raises worst case	Worst-case triggers (several already visible at cutoff) are converging in the May 11–21 window; the exit-triggered resumption pathway may be more dangerous than no-exit (KJ-1).

Owner	Indicator (open-source observable)	Signals	If observed, revise
	kinetic event following any U.S. unilateral victory declaration.		
MODERATE: periodic review; slower-moving structural signals			
T2	Provincial-vs-Tehran IRGC messaging asymmetry; absence of any documented Tehran recall of an in-progress provincial operation under pressure; concurrent provincial naval contacts in distinct sectors without unified signaling.	Supports	Accumulating support for the distributed-organizational actor model over the single-principal model; strengthens the case for the KJ-2 decentralization reading and the internal-spoiler scenario family.
T2 / T4	Hotel occupancy and aviation seat-load data for UAE and Qatar; source-country embassy advisory status; expatriate departure patterns within 24–72 hours of a symbolic civilian strike.	Confirms human-layer	Movement confirms the expatriate-workforce vulnerability that the physical-infrastructure tracker cannot capture; revise human-layer resilience and recovery-speed estimates (KJ-7).
T2 / T6	Repair-progress velocity at struck facilities (Fujairah F1 in particular); insurance reinstatement-or-withdrawal signals; nonlinear extension of repair timelines.	Tests compound pathway	Slower-than-modeled repair confirms the compound multi-layer failure hypothesis; recovery-preparation judgments (KJ-7) should be downgraded accordingly.
T7	MBS: a 24-hour gap in Royal Court communiqués or a missed public appearance; directives without MBS attribution; GCC statements missing Saudi-specific language.	Early warning	Early indicators of a Saudi leadership-incapacitation wildcard, a structural vulnerability independent of the war. Watch for the definitive trigger: confirmed incapacitation with no named successor within 48 hours.
T5	Unusual vessel concentration over the Omani-waters cable corridor at the mouth of Hormuz; GCC procurement of analog emergency communications or emergency water-rationing/military pre-positioning near labor concentrations.	Active planning	Indicates the EMP/cable-severance scenario is in active planning rather than assessment; elevate the T5 wildcard and pre-stage the analog-suppression tripwire (KJ-8).

Two features of this watch-list deserve emphasis. First, several of the highest-consequence indicators were already partially visible at the analysis cutoff. Mainly, mines being laid during the ceasefire, Hezbollah escalation signals, and suspended war-risk insurance, which is precisely why the worst-case end of the distribution was assessed as structurally closer than the best case in more than one domain. Second, the tiers are coupled: a Critical-tier confirmation rarely arrives alone, and the appearance of several High- or Moderate-tier signals in the same window should be read as a single converging alert rather than as independent data points. The instrument is therefore most useful not as a checklist but as a pattern detector; its purpose is to make a coordinated deterioration visible before any single indicator would, on its own, compel attention.

PART II: Domain Assessment

The seven chapters that follow are the evidence base on which Part I rests. Each is the work of one team examining a single domain of the crisis, and each was reviewed, red-teamed, and revised under the shared protocol described in the Methodology. They are presented here in a deliberate sequence rather than in team-number order. The assessment begins with the adversary whose behavior drives everything else (strategic intelligence), moves to the medium through which that behavior is contested (the maritime domain) and the systems it places at risk (critical infrastructure, then the economy), turns to the contest over perception that runs alongside the kinetic conflict (cognitive and information security), and closes with the two political questions on which the Gulf's future turns, the posture of its pivotal hedging state (the United Arab Emirates) and the durability of the collective itself (GCC regional stability).

Each chapter follows the same internal structure, situation and baseline, best and worst case, the critical what-if, the wildcard, and decision implications, and closes by handing its early-warning indicators forward to the consolidated watch-list in Part I. This common architecture is deliberate and allows the chapters to be read against one another as well as in sequence. Consistent with the information-cutoff convention stated in Part I, each assessment is anchored to conditions as of 17 April 2026; where a development after that date bears on a chapter's judgment, it is noted as a dated editorial update. Authorship is credited at the head of each chapter.

Iranian Intent and Command Coherence Under Decapitation

Strategic Intelligence Assessment Team (Team 1) · Jenn Huff, PhD

Situation and Baseline

The chapter's organizing question is deceptively simple: does Iran possess a coherent, centrally directed strategic intent, or is its wartime behavior driven substantially by decentralized provincial autonomy that persists regardless of what the center decides? The distinction is not academic. If tactical operations are governed by pre-authorized provincial commands, then progress at the political negotiating table need not produce any corresponding change in behavior on the ground, and the entire diplomatic architecture, which assumes Tehran can deliver de-escalation, rests on a mistaken model of the adversary. The assessment's central finding is that the second reading is the better fit. Iranian conduct is best explained not by centralized direction but by a thirty-one-province "Mosaic Defense" architecture in which each Islamic Revolutionary Guard Corps provincial command independently holds the weapons, intelligence, logistics, and authority to continue operating without instruction from Tehran.

The evidence for that finding is behavioral. During the ceasefire window, Iranian forces continued laying mines in the Strait of Hormuz, conducted fiber-optic drone strikes against targets in Kuwait and Iraq, and sustained a high tempo of allied-militia drone attacks on Saudi Arabia, all while Iranian diplomats were actively engaged in negotiation. That combination is difficult to reconcile with centrally directed restraint during a negotiating pause. It is readily explained by autonomous provincial action, a reading reinforced by an Iranian official's own attribution of early strikes to units that "could not be directly reached." The baseline that follows from this is not a forecast of resolution but of continuation. It is a managed stalemate that functions as a pressure-accumulation system, in which economic coercion, military reconstitution, domestic political pressure, and alliance dynamics compound until one of them forces a discontinuous change. The stalemate is the most probable near-term condition precisely because the structures sustaining it are stable.

Best and Worst Case

The best case is a negotiated agreement in which Iran provides verifiable commitments on the disposition of its highly enriched uranium in exchange for the lifting of the naval blockade and a phased normalization of Hormuz traffic. The team assessed this as remote, on the order of a few percent, because each of its enabling conditions is individually unlikely, and they must occur together. The Iranian leadership would have to authorize a verified transfer or third-party custody of the HEU; an external patron, Beijing or Moscow, would have to apply genuine compellence it has so far withheld; and a credible verification mechanism, possibly brokered by Oman, would have to be established. The structure of the best case is that it requires several low-probability events to coincide, which is why it sits far below the baseline.

The worst case is a resumed United States–Israeli air campaign converging with Hezbollah escalation, Houthi activation at the Bab el-Mandeb, and an Iranian sprint toward the nuclear threshold. The team rated this unlikely in absolute terms but assessed that the risk distribution is weighted toward this end rather than the middle, because several of its triggers were already visible at the time of writing. Whether Hezbollah strike activity, confirmed mine-laying, or explicit Iranian threats of "unprecedented" action if the blockade continued. Two features sharpen the judgment. The first is a near-term convergence window in mid-to-late May, when several trigger timelines overlap, with Russian reluctance to see the conflict reignite functioning as the principal structural dampener. The second is the most uncomfortable element of the analysis: an exit-triggered resumption, in which a unilateral declaration of victory and withdrawal invites an Iranian "parting shot," leaving the withdrawing party to choose between absorbing humiliation and resuming hostilities under worse conditions than before. The exit-triggered path may be more dangerous than no exit at all.

The Critical What-If

The chapter's counterfactual reverses the comfortable assumption that a United States withdrawal would de-escalate the conflict. Suppose instead that withdrawal takes the form of a unilateral victory declaration. The team traces this to what it calls the nuclear proliferation paradox: with the coalition disengaging, Iran retains its stockpile of roughly 440 kilograms of 60%-enriched uranium and faces every incentive to sprint to weapons-grade enrichment to foreclose any future return. The declared "victory" thus produces the worst available proliferation outcome while technically avoiding the worst kinetic one. The analytic value of the what-if is that it inverts an intuition. A decision-maker focused on ending hostilities without renewed fighting would read such an exit as success, whereas the structured analysis reveals that it locks in a nuclear-capable, Hormuz-sovereign, externally backed Iranian regional position that no later administration could address except at far higher cost. The forward inference is that the manner of any U.S. exit, not merely the fact of it, is a maximum-leverage variable.

The Wildcard

The wildcard formalizes the baseline's central insight as a low-probability, high-impact proposition: that Iranian command-and-control persists and continues executing operations regardless of central coherence, because the Mosaic Defense architecture makes leadership targeting far less decisive than standard doctrine assumes. The scenario is anchored in fact, the air campaign degraded a large share of Iran's missile stockpile and launchers, yet sub-threshold operations continued through the ceasefire, and in doctrine. The mosaic concept was articulated within the IRGC in the mid-2000s and restructured later in the decade, specifically to avoid the rapid collapse that befell Iraq's centralized command in 2003. The implication reaches beyond this conflict. If an adversary has deliberately engineered its forces for operational autonomy, then the analytic frameworks that predict command paralysis

following decapitation will systematically overestimate the effect of leadership targeting, a generalizable warning about mirror imaging that the other domain chapters inherit.

Decision Implications

The strategic recommendation is collective rather than bilateral. The Gulf states are converging on the recognition that they must build the security architecture that will outlast the conflict, regardless of how it ends. They are building it competitively, each hedging individually, which is precisely the fragmentation the Mosaic Defense doctrine was designed to exploit. The chapter frames the alternative as Saudi–Emirati co-convening of a collective framework, in which the states that jointly construct post-conflict Hormuz governance become its architects. While those who instead pursue separate accommodations with Tehran become individually dependent on Iranian goodwill, without the collective weight to negotiate from strength. The “scramble for Tehran” that Iranian messaging has sought to accelerate is, on this reading, a trap that competitive hedging walks directly into.

The Strait of Hormuz Under Dual Interdiction

Maritime Domain Team (Team 3) · Katie Hammond

Situation and Baseline

The maritime chapter asks whether Iran's stratified Hormuz access regime is evolving toward a multi-patron Gulf order, and what that shift means for GCC posture. The baseline judgment is precise: the Strait functions not as a closed chokepoint but as a politically filtered access regime. Daily transits had recovered from a low of two to four vessels to roughly twenty-four per day, against a pre-war baseline of over one hundred. This roughly 90% reduction occurs under conditions of dual interdiction, with Iran applying selective passage from the north and the United States blockading Iranian ports from the south. Iran permits transit for vessels tied to India, China, Turkey, and Pakistan while denying U.S.-allied shipping; an Oman-mediated proposal has begun to shift the access question from "which flags" to "which geographic lane." War-risk insurance remains suspended, lagging the recovery in physical transit. The most fragile of the baseline's assumptions is that Iranian and U.S. enforcement continue to operate in parallel without a deconfliction incident at sea.

Red team review reframed two elements. The team had treated the Oman proposal as a settled "controlled reopening" on single-source reporting; reviewers caught the anchoring and forced three competing hypotheses, controlled reopening, geographic stratification, and tactical Iranian feint, to be run in parallel. And the assumption that Iranian naval command operates as a unitary actor was identified as mirror imaging, echoing the strategic-intelligence chapter's central finding from the water.

Best and Worst Case

The boundaries are asymmetric, and that asymmetry is the chapter's key analytic move. The best case, a durable multilateral framework resting on Omani brokerage and a reconvening of U.S.–Iran talks, rose over the week from roughly ten to thirty-five percent probability on the strength of the Oman proposal and accelerating diplomacy. Despite the positive mood, it depends on three sequenced conditions holding simultaneously, none of which the GCC controls. The worst case, collapse of stratification into universal denial after a vetted vessel is struck, with war-risk coverage becoming unavailable and the market treating the Strait as closed at any price. Despite the lower headline probability it carries, it was judged structurally closer because it requires only one of several already-visible triggers to fire. The most uncomfortable of these is an accidental mine strike in the Omani-side lane. Iran has admitted it cannot account for all its mines. The Omani lane is exactly the space the reopening framework would activate, and a single such strike would discredit the geographic framework and force the deconfliction event neither navy has prepared for, the lowest-effort, highest-impact failure mode in the architecture.

The Critical What-If

The reversed assumption is that the conflict stays bilateral. The trigger is a visibly successful, casualty-producing Iranian strike on a U.S. carrier strike group using a Chinese-supplied weapon. This is the only single-variable change that converts a U.S.–Iran problem into a U.S.–China crisis, with the GCC as the operating theater, and the ingredients are already in place. The causal pathway runs in three phases:

1. Confirmed Chinese-weapon attribution.
2. U.S. escalation against the Iranian naval command and a public demand that Beijing renounce all assistance, with commercial shipping ceasing as insurance is withdrawn across all flags.
3. The Strait is fully shut, GCC basing drawn into kinetic exchange, and a Chinese carrier group is deploying to the Gulf of Oman to protect Chinese nationals.

The point is that all four pillars of the baseline collapse simultaneously rather than in sequence. Material weapon support, even one system, would convert China from a Hormuz transit beneficiary into a co-belligerent in U.S. perception, and GCC neutrality cannot survive that reframing.

The Wildcard

The wildcard is a formal Chinese naval escort declaration for Chinese-flagged LNG carriers transiting Hormuz, mirroring the institutionalized Gulf of Aden escort template extended westward. Red teaming made the wildcard more plausible, not less. The team had risked dismissing it as a black swan, but reviewers reclassified it as a known wildcard with the institutional capability already in place, dozens of PLAN escort task forces to the Gulf of Aden since 2008, an acute motivation in the share of Chinese LNG that transits Hormuz, and existing access infrastructure at Salalah and Djibouti. The failure mode is therefore attention, not imagination: the risk is missing the active indicators rather than failing to conceive of the event. If declared and sustained, the escort mission would be the most consequential change to Gulf maritime order since the establishment of the U.S. Fifth Fleet, the permanent end of single-patron Gulf security, after which every GCC decision on ports, escorts, and basing must be renegotiated continuously across two-armed patrons.

Decision Implications

The recommendation is to build the coordinating scaffolding before the crisis that proves it necessary: a GCC Maritime Coordination Cell, hosted in Oman and focused on deconfliction, incident reporting, and intelligence fusion, at an annual cost the chapter estimates at below $50 million. The cost of inaction is structurally higher. In a great-power crisis, GCC states will face strong incentives to underbid one another for Chinese or U.S. favor, and the absence of a coordination mechanism will produce sequential bilateral concessions that aggregate into a strategic position none of the six would have chosen. The overarching insight is invariant

across every scenario the team modeled: the decisive maritime change is not who transits Hormuz today but the permanent loss of Gulf maritime exclusivity, as the region moves from a single-patron to a contested multi-patron order.

The Energy–Water Nexus and Its Human Operating Layer

Critical Infrastructure Protection Team (Team 2) · Manny Ozigi / Brandon Woodie

Situation and Baseline

The infrastructure chapter's central methodological claim is that GCC critical-infrastructure protection cannot be modeled as a single-layer problem. Physical assets, the human workforce that operates them, and the diplomatic environment each have distinct failure velocities and recovery horizons, and configurations that fire on several layers at once exceed planning assumptions built for any one in isolation. The baseline is a middling outcome, more probable than either full stabilization or immediate collapse, but red-teaming made it darker and more prolonged rather than overturning it. The key surprise was that the controlling constraint on export recovery may be commercial confidence rather than physical passage. Insurance withdrawal and mine uncertainty can suspend commercially meaningful transit even when ships can technically move, shifting the analytic center of gravity from military access to system-level recovery conditions.

The chapter's most distinctive contribution is its treatment of the human layer. GCC infrastructure does not run on citizen labor; it runs on a predominantly expatriate workforce, and the baseline's standard model treats that workforce as a constant because Western infrastructure analysis can afford to. That assumption is the one the chapter's what-if is built to break.

Best and Worst Case

The best case is a constrained stabilization in which localized damage does not become a regional cascade. High-value coastal power-and-desalination complexes, such as Fujairah, are repaired quickly and not effectively struck. Insurers and shipowners treat movement as risky but manageable, and emergency measures buy enough time to prevent severe humanitarian or governance effects. It requires several enabling conditions to hold together, each of which is individually uncertain. The worst case is a sustained-campaign scenario in which repeated strikes and commercial lockout invalidate normal recovery assumptions across multiple GCC systems at once, the compound, multi-layer failure that the single-layer model cannot generate. Red-teaming's sharpest correction was that partial repair is not self-validating: repaired infrastructure can remain functionally unusable if crews, operators, or insurers still judge the threat environment intolerable.

The Critical What-If

The reversed assumption is the constancy of the expatriate workforce. The chapter's what-if surfaces a failure mode that the baseline structurally cannot produce. A strike that does not register on any infrastructure-damage tracker because it targets people, confidence, or working conditions rather than steel can nonetheless collapse what the infrastructure tracker exists to protect. A symbolic-reciprocal civilian strike that triggers expatriate departure

disables a plant as surely as a missile, and faster than repair crews can respond. The forward inference is that human-layer indicators, embassy advisories, departure patterns, occupancy, and seat-load data belong in infrastructure monitoring, not beside it.

The Wildcard

The wildcard is a single GCC-directed strike executed by a semi-autonomous Iranian actor, most plausibly a provincial IRGC command, outside Tehran's commitment authority and during a stated central de-escalation, with Tehran's subsequent disavowal reversing neither the strike nor its cascade. Its analytic value lies not in its probability, but in the failure mode it exposes: the GCC's diplomatic instruments lean heavily on Tehran as a responsible interlocutor, yet the Iranian side with which those instruments negotiate is not unitary. A strike from inside that fragmented system produces a sharp spike in coordination cost and escalation risk that the architecture is not optimized to absorb. Adversarial scrutiny narrowed the evidence base without falsifying it: the provincial-autonomy assumption remains the single highest-risk unfalsified assumption across the chapter's cascade architecture, the same proposition the strategic-intelligence chapter advances, here stress-tested against the infrastructure it would threaten.

Decision Implications

The recommendation is directed at GCC infrastructure operators and the GCC Secretariat's emergency coordination cell. Add an "internal spoiler / partial control" scenario family to the standard red-cell rotation, requiring multi-track signaling under attribution ambiguity, pre-positioned liaison capacity with labor source countries, emergency labor-retention authority, and air-defense allocation models that do not assume centrally coordinated Iranian targeting. The overarching takeaway is directional rather than probabilistic: the GCC is converging on a configuration in which its adversary is becoming less predictable while its own infrastructure becomes more brittle in specific, identifiable ways, and the two trends reinforce each other.

Hydrocarbons, Tourism, and the Speed of Recovery

Economic and Supply Chain Team (Team 4) · Angela Cummins / Kevin Izquierdo

Situation and Baseline

The economic chapter assesses the conflict across three pillars, hydrocarbons, tourism, and hospitality, and finds them more tightly coupled than they first appear. The baseline is a gray-zone recovery. It is elevated but not catastrophic, with Brent surging from around $72 to above $113, war-risk insurance suspended, and aviation and tourism operating well below pre-conflict capacity, yet recovery infrastructure intact and able to scale if security conditions hold. Red-team review judged the baseline analytically defensible but possibly too linear and too uniformly pessimistic about state-led recovery mechanisms, a useful counter to the chapter's own framing.

Best and Worst Case

The best case turns on a durable ceasefire formalized late in the second quarter, prompt mine clearance, and a phased reopening of Hormuz to escorted traffic. These actions allow producers to restore exports toward seventy to eighty percent of pre-conflict volume and Brent to ease toward the high eighties to ninety dollars by year-end. Travel advisories have been lifted, and Q4 is capturing a high-season rebound. The worst case is concentrated near the lower end of the range and is anchored by a specific trigger the team could not dismiss: a strike on desalination infrastructure, explicitly signaled in Iranian messaging. Because desalination is a force-multiplying target, disrupting residential stability, hotel operations, industrial output, and investor confidence at once, and triggering pressure to evacuate, it could occur even amid a limited escalation and would push recovery well beyond 2026, regardless of whether a formal ceasefire technically held.

The Critical What-If

The reverse assumption is that the Gulf's bypass routes provide genuine redundancy against a closure of the Strait of Hormuz. The trigger is a Houthi activation at the Bab el-Mandeb that closes the southern Red Sea. By that, sealing the last viable alternative export route and the Suez corridor simultaneously, without Iran firing a single additional missile, since Saudi tankers near Yanbu are already within confirmed Houthi range. The value of what-ifs is that they convert a single-chokepoint recovery problem into a network-vulnerability problem. The East–West pipeline and Red Sea access, which the baseline treats as partial buffers, are themselves geographically and militarily exposed. A simultaneous dual-chokepoint disruption would hit oil exports, shipping insurance, airline confidence, tourism, and hospitality at once. The most important risk is not physical closure alone but confidence collapse across multiple markets simultaneously.

The Wildcard

The wildcard is an Iranian nuclear detonation or credible demonstration, covert enrichment of the surviving stockpile of roughly 440 kilograms of 60% enriched uranium to weapons-grade, then a test or verified announcement. It is the one scenario in which every prior analysis, ceasefire timelines, insurance normalization, tourism recovery, sovereign-wealth resilience, becomes instantly irrelevant, and its preconditions are documented rather than theoretical, with IAEA access denied since late February. Red-teaming's most useful move was to identify the more dangerous form. Not just an overt detonation but a sustained verification blackout combined with ambiguous evidence of weaponization, in which uncertainty itself becomes the destabilizing force, damaging markets, tourism, and confidence without Iran ever crossing the overt-test threshold.

Decision Implications

The chapter's overarching takeaway is its title proposition: in a conflict the GCC did not start and cannot end, the only sovereign variable within its control is the speed at which it prepares to recover. Every scenario converges on the same truth: the post-conflict economic position will be determined not by the terms of the ceasefire but by the decisions made before it. The concrete recommendation is that the heads of government collectively and publicly authorize, within thirty days, a jointly governed Energy and Economic Reconstruction Fund pre-capitalized at no less than fifty billion dollars from existing sovereign wealth. The announcement is itself the action. It signals to every insurer, contractor, airline, and investor that the Gulf has already decided it will recover, compressing recovery timelines, reducing the Brent risk premium, and foreclosing the worst-case sovereign-wealth-liquidation spiral by demonstrating that sovereign capital is being deployed strategically rather than consumed reactively.

Narrative Integrity and the Architecture of Belief

Cognitive and Information Security Team (Team 5) · Lilya Avagyan

Situation and Baseline

The information chapter's baseline is managed deterioration, but its causal account is counterintuitive: the deterioration is driven less by Iranian capability than by two internal features of the GCC's own information architecture. The first is a speed asymmetry, centralized approval processes that leave official messaging chronically slower than the adversarial content it must answer. The second is an intra-GCC coordination gap. The chapter's defining finding, which red-teaming produced by reversing the team's starting assumption, is the self-defeating nature of information control. Tight control eliminates the independent verification infrastructure that is a society's primary defense against manipulation, so populations with no trusted source outside official channels are more vulnerable, not less, when those channels are discredited. There is nowhere to turn but the encrypted networks where adversarial content circulates uncorrected. The GCC's greatest information-security strength and its greatest cognitive-security weakness are the same system.

Best and Worst Case

The best case is calibrated resilience under pressure, conflict-driven institutional integration that peacetime diplomacy failed to deliver, in which a high-profile synthetic-media incident forces the GCC states to coordinate rather than compete. In addition, activating regional cybersecurity cooperation operationally for the first time and allowing official narratives to crowd out adversarial content before sectarian and loyalty frames harden. It requires several conditions to hold together, beginning with genuine, not merely declaratory, intelligence-sharing trust between Riyadh, Doha, and the UAE. The worst case does not require Iran to win the information war; it requires only that the GCC's own system fails visibly enough that populations stop believing it. The feature the team could not dismiss was the collapse of the credibility of official communication from within, not from adversarial success but from the architecture's own inadequacy under real-conflict conditions.

The Critical What-If

The chapter's what-if produces the single most arresting finding in the volume: the GCC cannot revert to an analog baseline because none was ever built. Every other society that has suffered catastrophic infrastructure failure had some pre-digital foundation, agricultural capacity, analog communications, and distributed food production to fall back on. The same modernization drive that made the Gulf one of the world's most digitally advanced regions made analog redundancy invisible as a priority, leaving no fallback layer at precisely the moment one would determine whether the society passes through a major infrastructure attack intact. The combination of internet blockage and a physical attack on the digital

substrate therefore does not produce a severe-but-recoverable crisis; it produces a discontinuity with no built-in path back.

The Wildcard

The wildcard pairs an electromagnetic-pulse event with the severing of the undersea cables that converge in a narrow seabed corridor at the mouth of Hormuz. Red-teaming made it more precise and more dangerous. The operative kill mechanism is not the instantaneous destruction of consumer electronics but the E3-pulse destruction of large power transformers. The latter are the components controlling power generation and desalination, the least replaceable on any conflict-relevant timeline, with manufacture measured in twelve to twenty-four months and no emergency airlift capacity. Reviewers also identified asymmetric blindness as more dangerous than mutual blindness: if Iran retains its hardened communications while GCC and U.S. systems go dark, the information asymmetry itself becomes the weapon, independent of physical destruction. The most dangerous early signal is not an Iranian deployment but GCC governments themselves substituting physical control mechanisms for digital ones, a sign they have intelligence that their own architecture may fail.

Decision Implications

The overarching takeaway is that the GCC's greatest cognitive-security vulnerability lies in the information-control architecture it built to protect itself. The centralized messaging, legal suppression of independent verification, and pervasive surveillance have simultaneously created the conditions for adversarial operations to exploit. It removed the resilience mechanisms that would absorb a cognitive attack and produced a single point of failure at a civilizational scale that no adversary needed to engineer. The highest-value action available is to build, now and before the signal fires, the response protocol for analog and satellite-communication suppression that no GCC state currently holds, the one signal in the analysis that offers hours, not days, of warning.

The Sustainability of the Hedging Posture

UAE National Coordination Team (Team 6) · Nam Phuong Ngo and Olivia Martin

Situation and Baseline

The UAE chapter examines the Emirates' tolerated-but-unendorsed hedge as a posture under strain, assessed as of mid-April and projected roughly thirty days forward. Red-teaming reinforced the baseline's probability but overturned its interpretation: the team's original frame, "Iran wants to destroy the UAE", was recency-and-drama bias, and adversarial review reclassified the pattern as coercion rather than annihilation. Selective passage for neutral-flag vessels, calibrated targeting, and avoided red lines (desalination, Jebel Ali) are all consistent with extracting concessions, which makes diplomacy more valuable, not less. The baseline rests most heavily on three falsifiable assumptions, that the Iranian leadership retains enough authority to direct the provincial commands, that Hormuz remains functionally closed to Western-allied traffic, and that the private-sector backstop and sovereign wealth funds continue to anchor the fiscal position. The chapter's key surprise is that the economic worst case is not a tail risk: in one pillar, it is assessed as roughly as probable as the baseline itself, so contingency planning must match continuity planning in urgency.

Best and Worst Case

The best case, "negotiated degradation," at low confidence, has a back-channel produce a renewable humanitarian pause, no further strikes on UAE energy infrastructure, a partial escort corridor, Brent settling to the mid-nineties, and the UAE positioning itself as the designated post-war interlocutor. It requires five independently uncertain conditions to hold at once. The worst case, "cascading collapse," assigned meaningful probability at medium confidence, has post-deadline strikes trigger Iran's threatened "irreversible destruction" of Gulf infrastructure. Coordinated strikes on Fujairah and Jebel Ali, a full Gulf insurance exclusion, a cyberattack on the central bank, Saudi well-killing that permanently reduces global spare capacity, and Houthi activation at the Bab el-Mandeb, with Brent breaking $150 and a structural expatriate-relocation cascade beginning. Each trigger was already visible at the time of writing.

The Critical What-If

The what-if reverses the baseline's tacit treatment of Fujairah and desalination as protected red lines, working backward from an end state in which both are non-operational. It exposes the entire diplomatic and economic architecture as resting on a single unverifiable premise, provincial IRGC restraint toward a handful of nodes, and identifies a concrete point of no return. Once ADNOC declares force majeure, there is no intervention point between that declaration and cascade completion. A desalination strike runs the cascade through the population rather than the export ledger, freshwater rationing for millions, the largely expatriate population facing compulsory rather than voluntary departure, and a thirty-day conflict becoming a multi-year structural loss as rival hubs capture relocated business. Red-

teaming added one refinement that the strategic-intelligence chapter would recognize: decentralization cuts both ways, since a provincial commander who can escalate autonomously can also stand down autonomously, making partial, uneven ceasefire compliance the most likely diplomatic texture.

The Wildcard

The wildcard is the only scenario in the volume that begins not with a missile but in Washington: a U.S. Treasury 30-day ultimatum demanding that UAE banks cease processing Iranian-linked transactions or face secondary sanctions designation and loss of dollar-clearing and SWIFT access. It matters disproportionately because crude is settled in dollars, a clearing cutoff strikes the mechanism for selling oil itself, and there is no air-defense system for sanctions. It targets precisely what every other actor is currently tolerating, the hedge, and forces the one binary the hedge cannot survive: full alignment with Washington or a pivot toward a yuan-denominated architecture. Red-teaming made it more plausible, not less, on the logic that the dollar-settlement dependency hands Washington more leverage per unit of pressure and thereby lowers the bar for Treasury action.

Decision Implications

The overarching takeaway is the one-line proposition that anchors KJ-5 in Part I. Across every scenario, whether the UAE ends the war as a mediator or a casualty, it comes down to a single choice. Whether it protects both desalination and the Fujairah corridor, or sacrifices one to save the other. Neutrality is the Emirates' most valuable strategic asset, and neutrality cannot survive a water crisis. The recommended actions are to pair desalination with the Fujairah export corridor under a single air-defense envelope before interceptor allocations are hardened. To open an Oman back-channel to propose a mutual humanitarian red line on water infrastructure, and to prepare a structured sovereign-wealth drawdown framework in advance so that any drawdown is coordinated rather than forced by panicked markets. The decision the UAE makes alone this week, the chapter argues, becomes the template the entire bloc inherits.

Cohesion, the Qatar Variable, and the Post-Conflict Order

GCC Regional Stability Team (Team 7) · Kayla Jordan / Maria Mancuso

Situation and Baseline

The final chapter assesses the durability of GCC political cohesion across four pillars, diplomatic, economic and energy, cyber and information, and political coordination. Its baseline is a fragile equilibrium holding at roughly fifty-five to sixty percent confidence, but the chapter's central contribution is to identify what that equilibrium actually depends on. GCC unity, it argues, is not a background condition of the regional environment; it is an active, Qatar-controlled variable with a specific observable tripwire, and no pillar team had been monitoring it as a shared priority. Qatar's structural compulsion to maintain channels to Tehran makes it simultaneously the weakest link in GCC solidarity and its most valuable mediator, a duality that cannot be managed by pressure alone, as the limited effect of the 2017–2021 blockade demonstrated.

Red-teaming shifted the baseline's center of gravity toward deterioration: three of the five worst-case triggers were already observable, while the best case required simultaneous progress across all four pillars, so the equilibrium was reassessed as biased toward decline in the absence of active intervention. The team also could not dismiss the possibility that Iranian information operations were already targeting GCC political cohesion directly rather than public panic, which would make the information pillar the primary vehicle of fracture rather than a secondary pressure.

Best and Worst Case

The best case has Qatar publicly reaffirm GCC unity in explicit terms, and an Omani back-channel produce a humanitarian-pause framework within a week. Iran accepts indirect negotiation through a GCC-credible third party, lifting confidence toward seventy-five to eighty percent, but only if every enabling condition materializes within a roughly two-week window, and only if external partners accept GCC neutrality as legitimate rather than pressing for alliance-side declarations. The worst case is the mirror image, in which Qatar's delay or qualification on a joint position is the first observable trigger, and the four pillars degrade together. The asymmetry is familiar from the other chapters: the best case requires a sequence to succeed; the worst requires only one visible trigger to fire.

The Critical What-If

The chapter's what-if reverses the assumption of Qatari alignment, and the result is the clearest demonstration of the variable's weight: reversing Qatar's co-signature behavior collapses all four pillars at once. This is what licenses the chapter's headline recommendation, that Qatar's co-signature on GCC joint security statements be designated a Priority 1 shared monitoring indicator across all pillar teams, with a defined forty-eight-hour revision trigger. If Qatar delays co-signature beyond that window, the worst-case probability should be revised

immediately from 25% to 45%, and the diplomatic, energy, and coordination pillars updated simultaneously.

The Wildcard

The wildcard is the incapacitation of the Saudi Crown Prince through illness, injury, or political removal, the only disruption that degrades all four pillars without requiring any Iranian action. It matters because Saudi strategic coordination is personally mediated rather than institutionalized. The operational capability of Saudi institutions to continue functioning is not in question, but strategic decision-making collapses without a single recognized authority capable of overriding institutional inertia and committing the Kingdom to novel courses. Red-teaming made it more plausible by surfacing exactly this capability–authority conflation. Because the trigger can arise from natural causes or domestic dynamics entirely outside the conflict, it is a structural vulnerability that contingency planning has historically underweighted.

Decision Implications

The overarching takeaway is that the multilateral institution is more fragile than its formal architecture suggests, and that monitoring resources are most efficiently concentrated on a single high-leverage indicator rather than distributed across diffuse signals. The relevant actor is the U.S. Special Envoy for Gulf Affairs in coordination with the GCC Secretariat, and the recommended discipline is twofold. Treat Qatar's co-signature behavior as the Priority 1 tripwire described above, and sustain the diplomatic investment in Qatar's stake in GCC institutional outcomes that pressure alone cannot substitute for. The team specified its own falsification condition, it would know it was wrong if Qatar co-signed every GCC security statement within twenty-four hours for fourteen days without qualification, and Qatar Energy opened no bilateral Tehran channel outside GCC frameworks.

PART III: The Outlook

The seven assessments were anchored to conditions as of 17 April 2026, and this report is written from a vantage point some six weeks later. That interval has been instructive. The stalemate the teams identified as the modal outcome has held, the fragmentation they warned of has, if anything, accelerated, and the single post-cutoff development of the first order, the United Arab Emirates' withdrawal from OPEC, was the competitive-fragmentation thesis made literal. The teams' framework has, in other words, aged well, which is the strongest available reason to take its forward-looking judgments seriously. This closing section does not restate the assessment; it draws out what the assessment implies about the period ahead.

The Stalemate Is a Process, Not a Resting Point

The most important thing to understand about the managed stalemate is that it is not a stable equilibrium but a pressure-accumulation system. Economic coercion, military reconstitution, domestic political pressure, and alliance dynamics do not cancel out at rest; they compound, until one of them forces a discontinuous change. The stalemate is therefore best read not as a pause in the conflict but as the conflict continuing by other means, with the kinetic tempo lowered while the underlying pressures build. This has a direct consequence for how the period ahead should be monitored. The relevant question is not whether the current quiet holds on any given day, but which of the accumulating pressures reaches its threshold first, which is precisely what the consolidated watchlist in Part I is built to detect.

The assessments also converge on an uncomfortable corollary: the outcomes that appear to resolve the conflict may be worse than continuing the conflict. A declared victory that leaves Iran nuclear-capable and Hormuz-sovereign is analytically worse than the stalemate it ends; a resumed air campaign that exhausts coalition resources without settling the nuclear or maritime-governance questions is worse than the stalemate it interrupts. The danger in the months ahead is thus not only escalation but the false off-ramp, the settlement that relieves immediate pressure while locking in a structurally worse position. Decision-makers reading for good news should be most skeptical precisely when an exit appears clean.

The Variables Most Likely to Break It

If the stalemate ends, it will do so through one of a small number of identifiable channels, and the assessments are unusually specific about those channels. The first is the convergence of several pressure timelines in a narrow mid-to-late-May window. Iranian crude-storage exhaustion, a ceasefire expiry, the rejection of a revised proposal, and renewed coalition military deliberation, any one of which could serve as the discontinuity the pressure-accumulation system has been building toward. The second is an autonomous provincial action that the center neither ordered nor can reverse, which would break the assumption on which every diplomatic instrument depends. The third is a single strike on a foundational infrastructure node, desalination, the Fujairah corridor, and an export terminal that converts a recoverable disruption into a structural loss. The fourth is the political wildcard that requires

no Iranian action at all: a rupture in the Saudi–Emirati relationship, or the incapacitation of the Saudi leadership, either of which would degrade the collective from within. These are not equally likely, but they are the channels worth watching; diffuse anxiety about “escalation” is less useful than attention to these four specific mechanisms.

Two external actors function as the principal brakes. Russia has a direct structural interest in an indefinite stalemate and the most immediate lever to discourage a coalition’s resumption, which makes it, paradoxically, a stabilizing force in the near term for reasons unrelated to Gulf security. China’s calculus runs through Hormuz energy dependence, and its restraint holds only so long as the conflict stays bilateral; the maritime chapter’s what-if marks the precise point at which that restraint would convert into involvement. The stability of the coming weeks rests partly on the continued, self-interested forbearance of two powers whose interests are not the Gulf’s own.

What the Window Is For

The single most actionable conclusion the seven teams reached, and the one on which they most nearly agree, is that the only variable fully within the GCC’s control is the speed at which it prepares. The preparation must be made now, during the conflict, precisely because the conditions may not improve. In every scenario the teams modeled, the states that pre-positioned capacity retained agency, and the states that waited inherited consequences. The implication is that the present interval of relative quiet is not a reprieve to be waited out but a window to be used. To pre-capitalize reconstruction capacity before a ceasefire rather than after it, to pair the irreplaceable infrastructure nodes under common protection before interceptor allocations harden. To establish the maritime and economic coordination mechanisms before the crisis, which proves that they are needed. And to build the analog and verification fallbacks whose absence the information chapter identified as a civilizational rather than a technical gap. In each case, the recommendation is the same in form: act during the pressure-accumulation phase, because the discontinuity, when it comes, will move faster than any deliberative process can.

The Choice the Gulf Is Actually Making

Beneath the seven domains lies a single decision that the Gulf states are making now, mostly without naming it. The war has not disrupted the regional security order so much as revealed that it had already ended, and the GCC states are not waiting for the conflict to resolve before constructing its replacement. They are building it in real time, using the war as both justification and forcing function. The only open question is whether they build it together or apart. Built collectively, the post-conflict architecture gives the Gulf the weight to govern its own waters, price its own risk, and negotiate from strength with the great powers now contesting its neighborhood. Built competitively, each state hedging individually, underbidding the others for external favor, racing one another to accommodate Tehran. This state produces exactly the fragmentation that the adversary’s defensive doctrine was

designed to exploit and leaves each member individually dependent on goodwill that none can guarantee. That is the choice the period ahead will force, and it is the one the Gulf still has time to make deliberately rather than inherit by default.

A Closing Note on the Method

A final observation concerns the exercise itself. The value of what these seven teams produced lies less in the accuracy of any single probability estimate than in the discipline the method imposed. Surfacing the assumptions beneath each judgment, tracing each alternative future to an observable indicator, and converting analysis into a monitoring program a decision-maker can actually run. That the framework has tracked the six weeks since the cutoff, the persistence of the stalemate, the acceleration of fragmentation, the literal arrival of the OPEC withdrawal it had anticipated in spirit, is a measure of the method more than of any forecaster's foresight. The conflict will continue to evolve in ways no assessment fixed at a single date can fully anticipate. What this report offers is not a prediction of where it ends, but a structured way of watching where it goes and, for the decision-makers who must act amid uncertainty, an argument for using the time that remains.



24 Minoos Str., Strovolos, Nicosia,
2042 Cyprus
info@strategyinternational.org
https://strategyinternational.org/