Twenty Years of Centralized Remedial Action at Southern California Edison: Architecture, Operational Lessons, and the Emerging Grid Control Platform

Bill Zhang, Obinna Obah, Ronald Lim, Steve Guerrero — Southern California Edison, Rosemead, California, USA

Joel Greene — SISCO, Inc., Stirling Heights, MI, USA

ronald.lim@sce.com

joel@sisconet.com

ABSTRACT

California's decarbonization mandate and Renewable Portfolio Standards (RPS) are driving rapid growth in renewable energy customers connecting to the Southern California Edison (SCE) transmission system. Remote wind and solar generation connects to the greater Los Angeles basin through a limited number of long-distance transmission corridors, creating contingency profiles that legacy stand-alone Remedial Action Schemes (RAS) were never designed to manage. As SCE's portfolio of stand-alone RAS grew to over 20 schemes—each independently engineered, hardwired, and commissioned—interactions between schemes became increasingly complex, commissioning timelines lengthened, and the maintenance burden scaled faster than available resources.

In 2006, SCE initiated Proof of Concept for a Centralized Remedial Action Scheme (C-RAS): a single, software-defined system locating all protection logic in the Grid Control Center and Alternate Grid Control Center, with IEC 61850 GOOSE communication to monitoring and mitigation relays distributed across the transmission network. This paper presents a practitioner account of what was built, what two decades of production operation revealed, and where the platform is heading.

The C-RAS architecture centers on a redundant Central Controller System comprising three Unified Analytic Platform (UAP) nodes per subsystem, implementing a two-out-of-three voting scheme to ensure correct mitigation actions are always taken. Today, the SCE C-RAS monitors and controls 17 substations, managing 362 IEDs across fully redundant communication paths, achieving a 50ms end-to-end fault-to-mitigation response time verified through independent Real Time Digital Simulator (RTDS) testing. This paper also presents the 2023 Layer 3 R-GOOSE migration as a replicable model for utilities operating wide-area IEC 61850-based systems and discusses the emerging role of C-RAS as a foundational platform for next-generation grid control applications in the DER and AI era.

1. INTRODUCTION

The electric transmission grid is being asked to do things it was never designed to do. The rapid integration of variable renewable energy—wind and solar—is displacing dispatchable generation with resources whose output is inherently uncertain, and whose location is often electrically remote from major load centers. In California, this transition has proceeded faster than in virtually any other jurisdiction in the world, driven by legislative mandates, market forces, and an abundance of renewable resources in geographically challenging locations.

Transmission planning and protection engineering have responded to this challenge in a variety of ways: new lines, enhanced monitoring, and updated operating procedures. But one of the most consequential responses—and one of the least visible to the broader industry—has been the development and deployment of Centralized Remedial Action Schemes. A C-RAS replaces a proliferating array of isolated, hardwired Special Protection Schemes with a single, software-defined system capable of wide-area coordination, dynamic reconfiguration, and real-time response across an entire transmission footprint.

Southern California Edison began this journey in 2006 and has now operated a production C-RAS continuously for more than a decade, through multiple technology generations, a major protocol migration, and a 2026 platform refresh. This paper, co-authored by the SCE engineering team responsible for the system and the technology vendor that has provided the central controller software throughout, presents a candid practitioner account: what was built, what the experience of operating it at scale has taught, and what the centralized architecture makes possible as the grid enters the next chapter.

2. THE TRANSMISSION CHALLENGE: CALIFORNIA'S RENEWABLE INTEGRATION PRESSURE

The deregulation of California's electricity market in 1996, combined with aggressive state-mandated renewable energy standards, produced a surge in independently owned generating plants connecting to the SCE grid. These plants—primarily wind turbines and solar photovoltaic facilities—are typically located in remote areas where open space and resource availability are abundant but transmission infrastructure is sparse. The power generated must travel to the greater Los Angeles basin load center over a limited number of long-distance transmission corridors.

As generation capacity on these corridors grew, power flows became increasingly difficult to predict and control. Under normal conditions, flows can approach corridor capacity. Under contingency conditions—the loss of a line, a transformer, or a generating unit—flows on remaining equipment can exceed thermal limits within seconds. The consequences include equipment damage, cascading outages, and potential for wide-area disturbances.

Mitigation options are limited. Building new transmission lines is expensive, time-consuming, and increasingly constrained by permitting and right-of-way challenges. Remedial Action Schemes offer a faster and more cost-effective path: automated detection of contingency conditions followed by rapid corrective action—generation tripping, load shedding, or system reconfiguration—before equipment limits are exceeded or stability is lost.

The three categories of power system problems that RAS must address in SCE's network are thermal overloads, voltage stability, and transient stability. Thermal overloads occur when contingency conditions force excess current through remaining equipment. Voltage instability can arise in heavily loaded systems with insufficient reactive power support, particularly on long transmission paths. Transient stability threats emerge when large disturbances—faults, sudden generation loss—challenge the ability of the remaining generation fleet to maintain synchronism. In each case, automated action within the sub-second window available is the only effective defense.

3. STAND-ALONE RAS: A SOLUTION REACHING ITS LIMITS

SCE had implemented more than 20 Remedial Action Schemes across its service territory since first deploying RAS technology. Each stand-alone scheme was designed as an independent system, with its own set of monitoring relays, trip relays, communication circuits, and control logic embedded in field equipment at substations. When a monitored contingency occurs, the scheme detects it locally and issues the programmed mitigation action—typically generator tripping or load shedding—within its isolated domain.

This architecture was appropriate when SCE had a small number of widely separated RAS schemes with minimal interaction between them. As renewable interconnection drove the proliferation of RAS applications, however, the limitations of the stand-alone approach became progressively more acute.

The fundamental problem is one of quadratic scaling. Each new RAS scheme does not merely add linear complexity—it adds potential interactions with every existing scheme. Conditions that trigger one scheme may affect the operating state that another scheme's logic was designed around. Coordinating the logic of multiple schemes requires understanding each in the context of all others, a combinatorial problem that grows unmanageable as scheme count increases.

Practical consequences compound the theoretical problem. Each stand-alone scheme requires its own hardware: relays, communication equipment, power supplies, and terminal space at affected substations. Commissioning each new scheme requires factory acceptance testing of the standalone components followed by multiple site acceptance tests at affected substations. Changes to any scheme—required when the connected generation fleet, the transmission topology, or the operating criteria change—require engineering analysis, hardware modification, and retesting of the modified scheme and its interactions.

SCE's engineering team reached a conclusion that an increasing number of utilities with complex generation interconnection environments are now reaching: implementing more stand-alone RAS schemes in a timely, cost-effective, and operationally sound manner was no longer feasible. A centralized approach was required.

4. C-RAS: ARCHITECTURE AND SYSTEM DESIGN

The Centralized Remedial Action Scheme differs fundamentally from stand-alone RAS in where the intelligence resides. In C-RAS, all protection logic is centralized in software running at the Grid Control Center (GCC) and its Alternate Grid Control Center (AGCC). Field equipment at substations is limited to monitoring and mitigation relays—devices that report system state upstream and execute trip commands on receipt of a mitigation signal but contain no local decision-making logic.

4.1 Central Controller System

The Central Controller System (CCS) is the heart of the C-RAS architecture. It is located at the GCC and AGCC, with each performing independent analytic functions. Each CCS comprises three Unified Analytic Platform (UAP) nodes operating in parallel. The CCS continuously processes incoming data from field relays, evaluates contingency conditions against arming equations, determines required mitigation actions, and transmits trip commands to field relays.

The CCS performs the following functions on a continuous basis:

- Processes field measurements: line and transformer flows, load and generation levels, substation topology
- Evaluates contingency conditions and detects contingency triggers from monitoring relay reports
- Interfaces with external data sources including the Energy Management System (EMS) for system state
- Calculates required arming states and determines appropriate mitigation actions
- Transmits mitigation commands to field relays via high-priority R-GOOSE messages
- Verifies that mitigation actions have been executed by monitoring relay acknowledgment

4.2 Voting Logic and Redundancy

The three-UAP architecture per subsystem is not redundancy in the conventional sense -- it is a voting scheme designed to eliminate false trip operations while maintaining high availability. Each UAP independently runs the full protection logic and, when a mitigation action is warranted, sends a separate

trip signal to the target relay. The field relay implements a two-out-of-three vote: it will execute a trip only when at least two of the three UAPs are online, and at least two have transmitted a coincident trip signal.

This architecture provides two critical guarantees simultaneously: First, no single failure can cause a false trip, since a single trip signal cannot satisfy the two-out-of-three criterion. Second, no single failure prevents a legitimate trip from executing, since two remaining UAPs can still satisfy the criterion. The system is fully redundant at the subsystem level as well, with independent A and B subsystems, each with their own CCS and communication paths.

4.3 Monitoring and Mitigation Field Relays

CRAS monitoring relays are installed at substations to report line and transformer status and flow information to the central controller. Line and transformer status information—contingency declarations—is transmitted as high-priority R-GOOSE packets, ensuring delivery within the timing budget required for sub-50ms response. Load flow measurements, used for arming calculations, are transmitted as lower-priority packets that do not require the same latency constraints.

The monitoring relay line-open declaration logic requires independent confirmation: either that breakers are open, and no current is flowing, or that protective relays have issued a trip command. This two-path confirmation prevents false contingency declarations from spurious breaker status indications.

CRAS mitigation relays are installed at generation and load tripping locations. On receipt of a valid two-out-of-three trip command from the CCS, they execute the specified generation trip or load shed. For locations outside SCE territory—where NERC CIP constraints preclude installation of R-GOOSE-capable relays—dedicated communication circuits with point-to-point protocols extend the trip signal from SCE's mitigation relays to relays at the interconnecting customer's site.

5. THE UAP PLATFORM: TECHNICAL IMPLEMENTATION AND PERFORMANCE

The Unified Analytic Platform (UAP) is the software engine that implements the C-RAS Central Controller System. Purpose-built for wide-area Special Protection System applications, UAP provides the analytics processing, communication interfaces, simulation, and integration capabilities required to operate a production C-RAS at the scale of SCE's transmission network.

5.1 Analytics Engine

The core UAP analytics engine processes incoming GOOSE and R-GOOSE data streams, evaluates contingency logic, and produces arming and mitigation decisions. The engine achieves a 4ms analytics reaction time from data receipt to decision output. Combined with the communication latency budget, this places the end-to-end fault-detection-to-mitigation-action time at 50ms in production operation—a performance level verified through independent RTDS full-transmission-system simulation and confirmed in on-site field acceptance testing at SCE.

5.2 Auto-Configuration from CIM and CID Files

A significant operational advantage of the UAP implementation is its ability to generate complete system configurations automatically from IEC 61968/61970 Common Information Model (CIM) files and IEC 61850 CID (Configured IED Description) files. When a new generating unit, substation, or relay is added to the C-RAS scope, the engineering team imports the relevant CIM and CID data, and UAP generates the corresponding monitoring and communication configuration without manual database entry. This capability has materially reduced commissioning time for new RAS applications and is a key differentiator of the centralized software approach versus hardwired alternatives.

5.3 Input Conditioning

Production experience with large IED populations reveals that field data is imperfect. Bad data points, communication interruptions, and IED behavioral anomalies can introduce erroneous inputs that, if unfiltered, could produce incorrect protection decisions. UAP implements per-IED input conditioning: agreement checking across redundant inputs, deadband limits to filter noise, and configurable bad-data identification. These functions run transparently as part of normal analytics processing, improving the quality of the data on which protection decisions are based.

5.4 Historian Integration and Scenario Replay

Every event processed by the UAP analytics engine is recorded to a historian—either OSIsoft PI System or InStep eDNA, depending on site configuration. The historian integration provides a complete time-stamped record of system state, contingency triggers, arming states, and mitigation actions. UAP also implements a Scenario Player capability that can replay any recorded event through the analytics engine for post-event analysis, diagnostics, or engineering review—without requiring live system access. This has proven particularly valuable for regulatory documentation and for training operations staff on rare high-severity contingency scenarios.

5.5 GOOSE Blaster and Monitoring

The GOOSE Blaster and Monitor utilities provide simulation and verification capabilities within the UAP environment. The GOOSE Blaster can generate synthetic GOOSE traffic simulating specific contingency sequences, allowing engineering teams to test protection logic in a lab environment before deployment. These scenarios can be hand crafted, or they can be derived from packet captures of live system data. The GOOSE Monitor provides real-time visibility into GOOSE and R-GOOSE message exchange across the system, supporting both commissioning verification and ongoing operational monitoring.

5.6 EMS Integration

UAP integrates with the utility's Energy Management System through a standardized interface, providing operators at the GCC with full arm/disarm control and status visibility. Operator displays are generated automatically from the CIM/EMS model, ensuring that the operational view of the C-RAS is consistent with the EMS topology model. Arm and disarm commands issued from the EMS are processed by UAP and reflected in the C-RAS analytics state. Every control action, configuration change, and state transition is logged to a Syslog-based audit trail queryable by the operations team.

6. TWENTY YEARS AT SCE: A DEPLOYMENT HISTORY

SCE's C-RAS deployment spans twenty years and five distinct phases. It represents the longest-running and most extensively deployed production C-RAS in North America, and its history provides a unique window into the practical experience of operating a centralized wide-area protection system through multiple technology generations.

6.1 Proof of Concept (2006)

The proof-of-concept phase had one critical objective: demonstrate that IEC 61850 GOOSE communication could achieve the latency required for wide-area protection over the distances and communication infrastructure characteristic of SCE's transmission network. The POC validated GOOSE messaging over a 500-mile fiber and microwave path, achieving a one-way latency of 21ms. This result confirmed the technical viability of the centralized approach and provided the foundation for regulatory engagement with the Western Electricity Coordinating Council (WECC).

6.2 Pilot Project (2007)

The pilot project established the production architecture: a dual-redundant system installed in two control centers, implementing the two-out-of-three voting logic in field relays and demonstrating full arming and mitigation functionality. WECC conditional approval for operation was granted in 2008 following completion of the pilot acceptance testing. The GE, SISCO, Cisco, and InStep team that would deliver the full deployment was established during this phase.

6.3 Full Deployment (2012)

Full deployment extended the C-RAS across 3 substations in the SCE transmission network, integrating 16 IEDs across the redundant A and B communication paths. The project involved factory acceptance testing at GE's facility followed by three site acceptance tests at SCE locations. The network infrastructure was designed and delivered by Cisco; historian integration was provided by InStep eDNA. The UAP central controller software, communications interfaces, and auto-configuration tools were provided by SISCO.

6.4 System Acceptance (2014)

Full system acceptance was completed in 2014 following a period of parallel shadow operation during which C-RAS decisions were verified against expected outcomes without affecting live system operation. The acceptance process included independent verification by RTDS full-transmission-system simulation—the most rigorous available standard for wide-area protection scheme validation. Acceptance marked the transition of C-RAS from a deployment project to a live production system under ongoing operations management.

6.5 Additional growth (2015-2018)

Several more RAS were phased in to a total of 5 - bringing substation and IED counts to 17 and 362 respectively.

6.6 R-GOOSE Migration (2019–2023)

The R-GOOSE migration, which began in 2019 and was completed in 2023, is described in detail in the following section. At its conclusion, all SCE C-RAS applications had been transitioned from Layer 2 GOOSE to Layer 3 R-GOOSE, eliminating the VLAN-based WAN workarounds that had accumulated since the initial deployment and placing the system on a standards-based, scalable communication architecture.

6.7 Platform Refresh (2026)

The 2026 platform refresh represents the current generation of the system, incorporating updated hardware, and enhanced software capabilities - supporting the growing customer base.

7. L3 R-GOOSE MIGRATION: A CASE STUDY IN STANDARDS-BASED MODERNIZATION

The SCE GOOSE-to-R-GOOSE migration is presented here as a detailed case study because the challenges it addresses—and the solutions it demonstrates—are broadly applicable to any utility operating a wide-area IEC 61850-based system. The migration was not merely a technology upgrade; it was a resolution of a fundamental architectural mismatch between a protocol designed for substation LANs and the demands of wide-area WAN communication.

7.1 The Problem with Layer 2 GOOSE on a WAN

IEC 61850 GOOSE was designed as a multicast protocol for the local area network environment of a single substation. In that context, its broadcast-to-all-listeners behavior is a feature: every device on the substation LAN receives every GOOSE message and discards those not intended for it. The protocol's fast-retransmission behavior—messages resent at increasing intervals following a state change—ensures rapid delivery on the low-latency, low-congestion substation LAN.

Extending GOOSE to a wide-area network introduces two distinct categories of problem. The first is a performance and stability issue: GOOSE was not designed for the WAN. When a grid event occurs, multiple relays simultaneously enter their fast-retransmission sequence. On a shared WAN, this produces a 'GOOSE storm'—a burst of multicast traffic that can saturate network capacity and interfere with other communication applications. Since every device on the network must receive and examine every GOOSE message to determine if it is the intended recipient, the compute and bandwidth burden grows with the number of devices.

The second is a cost and risk issue arising from the workarounds required to make Layer 2 GOOSE function on a WAN. The standard approach is to segment the WAN into Virtual Local Area Networks (VLANs), confining GOOSE traffic to the set of devices that need to communicate. This approach works, but it creates a non-standard, customized network configuration that expands in complexity with every new RAS application. Each VLAN requires configuration on every network device it traverses. As the number of RAS applications grows, the VLAN topology becomes a significant operational risk: more complex, harder to troubleshoot, more prone to configuration errors, and more difficult to recover from failure.

SCE's experience confirmed these effects. The VLAN architecture that had been put in place to support the Layer 2 GOOSE deployment was stable but rigid. Adding new C-RAS applications required careful VLAN engineering across the affected network segments. Troubleshooting communication anomalies required expertise in both the IEC 61850 protocol and the custom network configuration—a skill combination that is not easily maintained in an operations organization.

7.2 R-GOOSE: The Standards-Based Alternative

Routable GOOSE (R-GOOSE), standardized in IEC 61850-8-1, addresses the WAN deployment problem by encapsulating GOOSE messages in UDP/IP packets. This transforms the communication model fundamentally: instead of multicast broadcast to all devices on a network segment, R-GOOSE sends addressed packets directly to the specific subscribers that have registered for each publisher's data. The result is targeted delivery that scales with the number of relevant subscribers, not with the total number of devices on the network.

The network architecture implications are equally significant. Because R-GOOSE operates at Layer 3, it can be routed using standard IP networking equipment and practices. No VLANs, no custom configuration, no specialized knowledge of GOOSE-specific network behavior is required. The network simply routes packets. This is the same model used for every other WAN protocol in the utility's infrastructure, meaning the skills and tools required to operate it are broadly available.

7.3 Migration Approach and Execution

The SCE migration from Layer 2 GOOSE to R-GOOSE involved four categories of work:

- Telecom network changes: modifications to add R-GOOSE connectivity to the existing substation communication infrastructure, establishing routed paths between the central controller locations and each substation
- Equipment replacement: replacement of Layer 2-only relays with R-GOOSE-capable relay models, executed as a phased rollout across the 6-substation footprint

- System monitoring: addition of monitoring capability for the new L3 C-RAS alongside removal of monitoring for the L2 system as substations were migrated
- Testing and commissioning: thorough verification of R-GOOSE communication performance, latency, and behavior at each migrated substation, confirming that the 50ms end-to-end performance specification was maintained

The migration was designed to be executed without disruption to C-RAS protection service. Substations were migrated in sequence, with each migrated substation verified on R-GOOSE before the L2 path was decommissioned. At no point was the overall C-RAS protection function unavailable.

7.4 Outcome

The migration was completed in 2023. All SCE C-RAS applications now operate on the L3 R-GOOSE protocol. The VLAN architecture has been retired. The operational benefits are concrete: the communication infrastructure is now maintained using standard IP networking tools and practices, troubleshooting is substantially simpler, and the addition of new C-RAS applications does not require custom network engineering. The migration also aligns the SCE C-RAS with the direction of the industry: R-GOOSE is the current standard for wide-area IEC 61850 deployments, and equipment and expertise are increasingly available accordingly.

8. OPERATIONAL LESSONS FROM TWENTY YEARS OF PRODUCTION C-RAS

Operating a production C-RAS for two decades across a major transmission network generates experience that cannot be obtained from laboratory testing, simulation, or analysis. The lessons presented here reflect the accumulated operational knowledge of SCE's protection engineering team and are offered as practical guidance for utilities evaluating or implementing centralized wide-area protection systems.

8.1 Stakeholder Coordination Is Non-Negotiable

A C-RAS touches more parts of the utility organization than any stand-alone RAS scheme. Protection engineering must coordinate with transmission planning (to ensure the scheme addresses the right contingency scenarios), substation engineering (for field relay installation and configuration), operations at the GCC (for operator interface and training), and telecommunications (for the communication infrastructure). In SCE's experience, the projects that proceeded most smoothly were those in which all stakeholder teams were engaged early and maintained throughout. Communication gaps between groups—particularly between protection engineering and telecom—were a source of delay and rework.

8.2 Redundancy Design Is Worth the Investment

The dual-subsystem, three-UAP voting architecture adds cost and complexity at the design stage. Two decades of operation have confirmed that this investment pays for itself many times over. Failures of individual components—relay modules, communication circuits, UAP nodes—have occurred over the operational life of the system. In every case, the redundancy architecture has prevented any such failure from affecting protection service. The two-out-of-three voting logic has also, in multiple instances, prevented spurious signals from producing incorrect operations. Utilities evaluating C-RAS architectures should regard full redundancy not as an optional enhancement but as a fundamental design requirement.

8.3 Centralized Logic Changes Are Dramatically Simpler

The most significant operational advantage of the centralized approach—and the one most difficult to appreciate before experiencing it—is the ease of modifying protection logic in response to changing grid conditions. When new generation interconnects, when the transmission topology changes, or when WECC operating criteria are revised, the corresponding logic changes can be implemented in software at the

central controller. No relay replacement, no substation rewiring, no site visits for field configuration changes. The change is engineered, simulated, and tested in the UAP environment, then deployed. This has allowed SCE to maintain an up-to-date protection scheme in a transmission environment that has changed substantially since initial deployment.

8.4 Regulatory Engagement Must Be Proactive

C-RAS operates under NERC reliability standard PRC-012-2, which establishes requirements for Remedial Action Schemes intended to prevent unintentional or unacceptable reliability risks. In addition, C-RAS applications in the Western Interconnection require review and approval by the WECC Remedial Action Scheme Reliability Subcommittee (RASRS) and the Reliability Coordinator. SCE's experience is that proactive engagement with these bodies—beginning well before formal submission and maintaining transparency about system design, testing methodology, and operational experience—substantially reduces the time and uncertainty of the approval process. The parallel shadow operation approach, in which the C-RAS evaluates contingencies and calculates actions without executing them, is a particularly effective tool for demonstrating performance to regulators before seeking operational approval.

8.5 Documentation and Audit Trail Are Operational Assets

The historian integration and full audit trail capabilities of the UAP platform were specified at the outset as engineering requirements. In practice, they have proven to be operational assets that support the entire system lifecycle. Post-event analysis following real contingency operations, regulatory inquiries and compliance audits, engineering reviews during scheme modifications, and training of new staff have all relied substantially on the documented event history. Utilities considering C-RAS implementations should ensure that historian integration and comprehensive logging are treated as first-class system requirements, not afterthoughts.

9. C-RAS AS A GRID CONTROL PLATFORM: BEYOND PROTECTION

The SCE engineering team has publicly articulated a vision that extends substantially beyond conventional protection: C-RAS as the foundational platform for next-generation grid control. This vision is grounded in a structural observation about what a production C-RAS already provides. A system that aggregates real-time status and measurement data from 120 substations, communicates bidirectionally with field devices at sub-50ms latency, executes complex coordinated logic centrally, and maintains a historian of every system event is, architecturally, a real-time grid control platform. The applications that can be implemented on that platform are limited primarily by imagination and regulatory framework, not by technical capability.

SCE has identified a number of specific applications for which the C-RAS architecture provides a natural implementation path:

9.1 Wildfire Risk Mitigation

Downed power lines are a significant ignition source for wildfires in California's fire-prone terrain. A broken conductor that remains energized and contacts the ground can ignite vegetation before conventional overcurrent protection operates. Synchrophasor measurement data, combined with the centralized analytics capability of the UAP platform, can detect the electrical signature of a broken conductor—high-impedance fault—and interrupt line energization before ground contact occurs. The C-RAS communication and control infrastructure provides a natural implementation layer for this application, which requires both fast, distributed sensing and coordinated response.

9.2 Automated Blackstart

Grid restoration following a major outage requires precise coordination of generation startup, transmission energization, and load pickup sequences. Today, this coordination is largely manual, requiring direct communication between control center operators across multiple organizations. The C-RAS, with its centralized logic and wide-area control capability, provides the foundation for automating blackstart sequences: executing programmatic restoration procedures that would otherwise require manual coordination, reducing restoration time and the risk of errors under high-stress conditions.

9.3 Grid-Scale Battery Dispatch

Distributed battery storage assets are increasingly connected across the SCE transmission and distribution network. Coordinating their dispatch for grid stability purposes—absorbing excess generation, providing fast frequency response, supporting voltage stability—requires real-time situational awareness and fast response capability. The UAP analytics layer, already receiving measurement data from across the C-RAS substation footprint, can evaluate dispatch decisions and communicate control signals to storage assets over the same R-GOOSE infrastructure used for protection.

9.4 Dynamic Line Ratings and Voltage Control

Conventional transmission line ratings are static, based on worst-case ambient conditions. Dynamic Line Ratings (DLR) use real-time temperature, wind, and line loading data to compute actual thermal capacity, which is often substantially higher than the static rating. The C-RAS, with its real-time measurement infrastructure, can incorporate DLR data into arming calculations—allowing higher utilization of existing transmission capacity and deferring expensive new construction. Similarly, the wide-area measurement and control infrastructure supports centralized Voltage and VAR optimization, extending UAP analytics to reactive power management across the transmission footprint.

10. THE AI AND DER INTEGRATION HORIZON

The C-RAS architecture has a property that becomes increasingly significant as artificial intelligence and distributed energy resource management emerge as grid operations priorities: it is inherently a centralized data aggregation and control system. Every AI application requires data; C-RAS continuously collects it from 17 substations. Every DER coordination application requires communication and control capability; C-RAS has it, with verified sub-50ms performance.

10.1 AI-Assisted Analytics

Bill Zhang, SCE's subject matter expert for RAS and a co-author of this paper, has identified several specific areas where AI can enhance C-RAS operation. Mitigation algorithm improvement using AI to identify anomalies and patterns in historical event data—recognizing precursor signatures of grid stress conditions before they escalate to contingencies—represents a natural extension of the existing analytics capability. Machine learning models trained on the historian record can provide early warning of developing conditions, complementing the deterministic protection logic with probabilistic situational awareness.

10.2 Engineering Design Automation

The design of C-RAS protection logic—arming equations, mitigation tables, contingency classifications—is currently a labor-intensive engineering task. AI-assisted design tools could substantially accelerate this process: analyzing the transmission topology and generation interconnection queue to propose initial scheme designs, running simulation cases to verify performance, and generating the relay settings and UAP configuration files required for implementation. The auto-configuration capability already present in UAP provides a natural integration point for such tools.

10.3 AI-Driven Simulation and Testing

The UAP Scenario Player capability, which allows replay of recorded events through the analytics engine, provides a foundation for AI-driven testing. Generative models could create synthetic contingency scenarios, including rare and high-consequence events that may not appear in the historical record—for testing against the protection logic. This would allow more comprehensive coverage of the contingency space than is achievable through deterministic scenario development alone, potentially identifying logic gaps before they are exposed by real events.

It is important to characterize these AI applications honestly. The deterministic protection functions of C-RAS—the arming logic, the voting scheme, the mitigation execution—are, and should remain, deterministic: predictable, verifiable, and regulatory-compliant. AI applications are most appropriately positioned as advisory layers that enhance situational awareness and accelerate engineering processes, not as replacements for the deterministic protection core. The regulatory framework for AI in protection applications is still developing, and practitioners should engage proactively with NERC, WECC, and their regional entities on how AI-assisted functions will be characterized and approved.

11. CONCLUSION

Southern California Edison's twenty-year experience with centralized remedial action systems demonstrates that the centralized approach to wide-area protection is operationally viable, technically superior to stand-alone schemes at scale, and increasingly foundational to grid control applications that go well beyond conventional protection.

The key conclusions for utilities evaluating or planning wide-area protection systems are:

- Centralized RAS is proven at scale. SCE has operated a production C-RAS continuously for over a decade, managing 17 substations and 362 IEDs with a 50ms end-to-end response time. This is not a laboratory result or a projection—it is documented, verified, production performance.
- The Layer 3 R-GOOSE migration is replicable and substantially reduces operational risk. Utilities operating Layer 2 GOOSE over WAN should evaluate migration to R-GOOSE not as a future option but as a near-term priority. The VLAN complexity and operational risk of the Layer 2 approach grow with every new application.
- The centralized architecture is a platform, not just a protection scheme. The same infrastructure that enables sub-50ms protection response—centralized analytics, wide-area measurement, bidirectional field communication—also enables wildfire mitigation, automated restoration, DER coordination, and AI-assisted operations. The utilities that invest in this architecture today are building the operational foundation for grid control capabilities that will be required in the DER era.

The authors acknowledge the contributions of the SCE and SISCO engineering teams who have designed, built, operated, and continuously improved this system over two decades, and of the WECC RASRS Subcommittee and Reliability Coordinator for their engagement in the regulatory approval process.

REFERENCES

[1] B. Zhang and S. Rodriguez, "Powering the Future: How SCE is Revolutionizing Grid Stability with Centralized Remedial Action Schemes (CRAS)," PAC World Magazine, Issue 072, June 2025.

[2] IEC 61850-8-1, "Communication networks and systems for power utility automation — Part 8-1: Specific communication service mapping (SCSM) — Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3," International Electrotechnical Commission.

[3] IEC 61968/61970, "Application integration at electric utilities — System interfaces for distribution management / Energy management system application program interface (EMS-API)," International Electrotechnical Commission.

[4] NERC Standard PRC-012-2, “Remedial Action Schemes,” North American Electric Reliability Corporation.

[5] WECC RASRS, “Remedial Action Scheme Reliability Subcommittee Guidelines,” Western Electricity Coordinating Council.

[6] B. Coalson, A. Ang, and M. Avendaño, Southern California Edison, “C-RAS as a Grid Control Platform,” IEEE Smart Grid eNewsletter, August 2020.

[7] IEEE Std C37.239-2010, “IEEE Standard for Common Format for Transient Data Exchange (COMTRADE) for Power Systems,” IEEE.

Note: Reference [6] details should be verified with SCE. Additional references to be added as appropriate for the full paper.