Electronic Design®

By ALAN R. EARLS, Contributing Editor

Attack Grid Power Reliability Issues with Control System Redundancy

Reduced reliability in grid power and other risks make it timely to consider building in greater resiliency. It's not for the faint-hearted, but with a step-by-step approach, it can be accomplished without breaking the bank.

Power outages are becoming more frequent in the U.S. One research group states that from 2011 to 2021, the U.S. experienced 64% more power outages than in the previous decade (2000-2010). Much of this increase has been attributed to more frequent and severe weather events.

Other factors contributing to the problem are that a significant portion of the U.S. power grid was built decades ago and it's aging, making it more susceptible to failure, plus there's increased demand. Growing populations, the electrification of vehicles and buildings, and the energy demands of artificial intelligence (AI) are putting more strain on the power grid, increasing the likelihood of outages.

Lastly, attempts to switch from large, centralized power plants to diverse, distributed, and often less-predictable power sources have triggered an, as-of-yet, incomplete rebuild of the grid and its support systems. How this will ultimately work out is unclear.

All of this is bad news for process-type industries, where outages can lead to significant disruptions, including halted production lines, material loss, supply-chain disruptions, and potential safety hazards. These can result in financial losses, damaged equipment, and even risks to employee safety.

TABLE 1. APPROACHES TO ACHIEVE PLC SYSTEM REDUNDANCY

Redundancy Type	Hardware Required	Switchover Time	Reliability	Cost
Full hardware	Full duplication	<1 second	Very High	High
Partial redundancy	Partial duplication	Several seconds	Medium	Medium
Software-based	Minimal (uses IT/ SCADA)	Varies (slower)	Low-Medium	Low

Mitigating the impact of unplanned power outages in industrial process-control systems is essential to maintaining safety, product quality, and system integrity.

What Steps Can Engineers Take to Help Mitigate Power Outage Impacts?

Many factors contribute to resiliency. A well-designed strategy for mitigating the impact of unplanned power outages could include a data-retention strategy, automated recovery, and failover planning, in addition to, of course, power protection such as generators or batteries.

One of the most effective approaches is to implement industrial-automation system redundancy. It can also be the most expensive way to maintain system availability during hardware or power failures, making it important to have a well-thought-out approach.

By providing for seamless failover to standby control systems, there should be no loss of process control or operator visibility. However, this requires investing in redundant programmable logic controllers (PLC) or programmable automation controllers (PAC), and/or implementing hot-standby controllers with redundant power supplies and field I/O modules. For good measure, figure in a dual Ethernet ring. [This may be beyond the means of most operations, but it could make sense with high-value-add products and/or particularly demanding setup and restart requirements.]

What are Some Paths to PLC Redundancy?

To make a process system

with a PLC redundant, you don't always have to duplicate the hardware exactly, but hardware duplication is the most straightforward and reliable method. Three main approaches can be taken to achieve redundancy in PLC systems, depending on the level of reliability, cost, and complexity you're willing to implement:

1. Full hardware redundancy (also called hot-standby or synchronous redundancy)

This involves deploying two identical PLCs (same brand, model, firmware) running in parallel. One is the primary, and the other is a hot standby. The key features of full hardware redundancy are real-time synchronization of memory and I/O and automatic switchover if the primary fails. The arrangement often includes redundant power supplies, network interfaces, and I/O modules as well.

- Pros:
- High availability and fast failover
- Minimal process disruption
- Cons:
- High cost due to duplication of hardware and licenses

2. Partial redundancy (shared or switched resources)

This involves duplicating only the controller, while some components (e.g., I/O racks or human-machine interface) are shared or connected via switchover logic. For example, this could be implemented with a single I/O that has dual PLCs and switchover relays or software-controlled switchover logic.

- Pros:
- Cost savings over full duplication
- Some improved fault tolerance
- Cons:
- More complex to manage
- Longer recovery time
- Shared points (like I/O) can still be a single point of failure

3. Software-based redundancy or high-level supervisory redundancy

This involves implementing redundancy logic in software, possibly in a supervisory control system (like SCADA or DCS) or higher-level control utilizing distributed computing. The approach uses heartbeat checks and watchdog timers or SCADA or external system monitors where the PLC initiates switchover.

- Pros:
- Flexible
- Can leverage existing IT infrastructure
- Cons:
- Less deterministic
- Dependent on network health and SCADA performance

The required redundancy level for PLCs depends on the criticality of your process. For high-stakes applications (e.g., chemical plants, power generation), full hardware redundancy is preferred. For less-critical systems or where budget is constrained, partial or software-based redundancy can still provide reasonable fault tolerance.

What are Some Paths to PAC Redundancy?

When designing redundancy for a PAC system, you don't always have to duplicate the hardware exactly. However, as with PLCs, the level and method of redundancy depend on the criticality of the system, performance requirements, and cost tolerance. PACs offer more flexibility and integration capabilities than PLCs, and thus have more nuanced redundancy strategies:

1. Full hardware redundancy (controller + I/O)

This is when you employ two identical PACs with duplicated I/O, power, and communication interfaces. Both controllers run the same program—one is primary, the other is synchronized standby; an approach favored in mission-critical systems (e.g., water treatment, pharma, power). Commercial examples include Allen-Bradley ControlLogix Redundancy and Siemens S7-400H.

- Pros:
- Fast failover (sub-second)
- High availability
- Cons:
- Highest cost due to full duplication

2. Partial hardware redundancy

This is a setup with redundant PACs but has shared or

TABLE 2: PAC OPTION COMPARISON

Redundancy Method	Hardware Duplication	Fast Failover	Cost	Typical Use Case
Full PAC + I/O redundancy	Yes (all)	Yes (<1 sec)	High	Critical infrastructure, pharma
Partial hardware redundancy	CPU only	Moderate	Medium	Industrial systems with cost limits
Network/distributed control	No (shared logic)	Moderate	Medium	Modular manufacturing
Virtual redundancy	No (virtualized)	Yes	Medium	Software-defined automation
SCADA/IT-level redundancy	No	No (slow)	Low	Data logging, batch/ SCADA override

multiplexed I/O modules. It employs a redundant CPU with shared I/O racks and switchover logic that selects which CPU is active. It's common in systems where I/O duplication is too expensive or space-limited.

- Pros:
- Balanced cost
- Resilience
- Cons:
- Shared I/O can be a single point of failure

3. Network redundancy + distributed control

This involves distributed PACs across redundant industrial Ethernet (e.g., EtherNet/IP, PROFINET) or ring networks (e.g., MRP, HSR, PRP). It works because control is distributed; one PAC failure won't halt system-wide functionality. The approach is built on redundant gateways, dual NICs, and a ring topology.

- Pros:
- Modular and scalable
- Higher fault isolation
- Cons:
- Complexity
- Not suitable for tight loop control redundancy

4. Virtual redundancy / PAC virtualization

This is when PAC functions are virtualized on industrial PCs or hypervisors with redundant failover. The PAC software (e.g., SoftLogix, Codesys) runs in virtual environments with VM-level failover. The approach is often applied in modern, hybrid control systems where the PAC is software-defined.

- Pros:
- Flexible and scalable
- Easy backups and snapshots
- Cons:
- Dependent on hypervisor/IT infrastructure stability

5. Supervisory / SCADA-level redundancy

In this case, redundancy is handled at the SCADA or MES layer rather than inside the PAC itself. The PACs may not be redundant, but the supervisory system can take over control logic or reroute around faults.

- Pros:
- Software-based
- Cheaper
- Cons:
- Slower response
- Riskier for fast control loops.

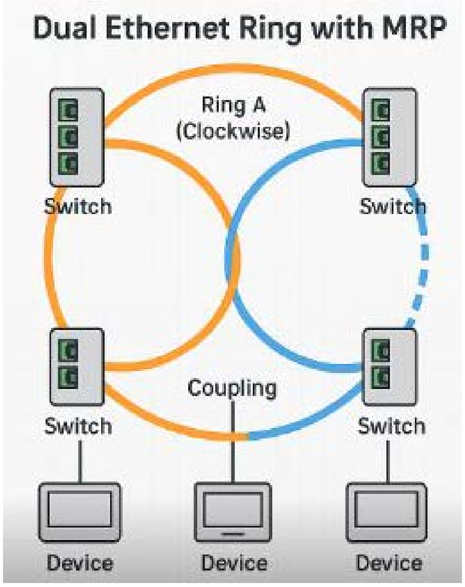

Shown is a generalized configuration redundant Ethernet ring employing the MRP (Media Redundancy Protocol) – IEC 62439-2, which is one of the options for ring topologies.

How Can a Dual Ethernet Ring Help with Resilience?

Network resilience matters, too. As an important component in many control systems, adding redundancy to an Ethernet ring can be a cost-effective investment. To build a dual Ethernet ring, you're creating a redundant, fault-tolerant network topology where two counter-rotating Ethernet rings provide resilience and fast failover in industrial or critical systems. This is common in industrial Ethernet, substation automation, and mission-critical infrastructure.

Start off by understanding the nature of your network and its important requirements, such as low recovery time (typically <50 ms), high reliability, and/or deterministic communication. Then, choose a ring redundancy protocol that supports dual ring setups.

Popular protocols include:

- **HSR (High-availability Seamless Redundancy):** IEC 62439-3; sends duplicate frames over both rings.
- **PRP (Parallel Redundancy Protocol):** IEC 62439-3; used with HSR or separately.
- **MSTP/RSTP with Ring Coupling:** Traditional spanning-tree-based redundancy.
- **MRP (Media Redundancy Protocol):** IEC 62439-2; more efficient than RSTP for ring topologies *(see figure)*.

You will also need managed industrial Ethernet switches with support for:

- Dual ring redundancy
- Fast failover (typically ≤20–50 ms)
- Protocol support (HSR, MRP, etc.)

Vendors include Hirschmann, Moxa, Siemens, Cisco (IE switches)

Network design requires creating two independent physical rings, namely a clockwise Ethernet ring and a counter-clockwise Ethernet ring. Then each device or switch is connected with two ports to each ring (dual-NIC or dual-port switch).

You can employ redundancy boxes (aka redboxes) if devices are single-attached. Redboxes can connect single-port devices to a dual-redundancy ring, or they can connect singly attached nodes (SANs) to the redundant network. Essentially, a redbox acts as a "proxy" for a node that doesn't have two network connections, enabling it to participate in the redundant network as if it had a dual connection.

In protocols like MRP, one switch acts as the ring manager to monitor topology changes and manage ring closure.

Redundancy involves configuring the chosen protocol on each switch. You may have to designate ring managers/supervisors as needed and validate settings like heartbeat timing, frame duplication, and hold-off times.

Finally, test for redundancy! For example, you can simulate cable or port failures to verify automatic failover between Ring A and Ring B and then monitor recovery time and network continuity.

A few other tips that can help make resilience a reality:

- Use fiber for long distances between switches.
- Provide a separate power supply for each switch.
- Monitor using SNMP or a network management system (NMS).
- Clearly label cables and document topology.
- Don't miss the potential value of retimer technology, particularly where long cable runs are present.

References

Demystifying redundancy in automation
Understanding US Power Outages