Distributed Systems
600.437
Intrusion-Tolerant SCADA

Department of Computer Science
The Johns Hopkins University

Network-Attack-Resilient Intrusion-Tolerant SCADA for the Power Grid

Lecture 10

Further readings:

- *Survivable SCADA via Intrusion-Tolerant Replication*, Jonathan Kirsch, Stuart Goose, Yair Amir, Dong Wei, Paul Skare, IEEE Smart Grid 2014.
- *Toward Survivable Intrusion-Tolerant Open-Source SCADA*, Thomas Tantillo, DSN Student Forum 2015.
- *Network-Attack-Resilient Intrusion-Tolerant SCADA Architecture*, Yair Amir, Amy Babay, Thomas Tantillo, U.S. Provisional Patent Application No. 62/353,256, 06/2016.

Importance of SCADA Systems

- Supervisory Control and Data Acquisition (SCADA) systems form the backbone of critical infrastructure services
  - Power grid, water supply, waste management
- To preserve control and monitoring capabilities, SCADA systems must be constantly available and run at their expected level of performance
- SCADA system failures and downtime can cause catastrophic consequences, such as equipment damage, blackouts, and human casualties

Modern SCADA System Architecture

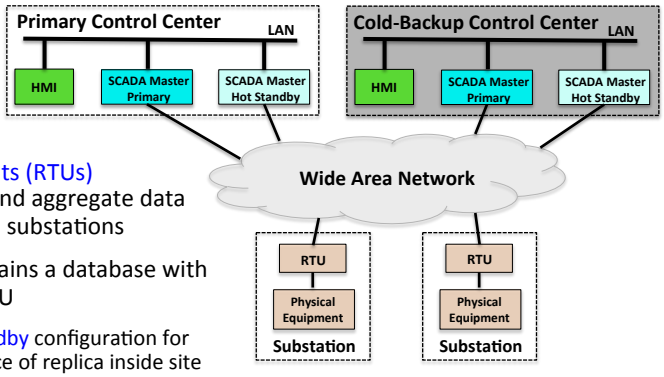

- Remote Terminal Units (RTUs) communicate with, and aggregate data from, local sensors in substations
- SCADA Master maintains a database with the status of each RTU
  - Primary / Hot Standby configuration for crash fault tolerance of replica inside site
  - Primary / Cold Backup configuration for crash fault tolerance of an entire site
- Human Machine Interface (HMI) provides graphical displays for operator
- Timeliness requirements of 100 – 200 milliseconds for critical monitoring and control data

SCADA Migrating to IP Networks

- Traditional SCADA systems ran on proprietary networks
  - Created air gap from outside world and attackers
- Cost benefits and ubiquity of IP networks are driving SCADA to use IP networks
  - Exposes SCADA to hostile environments, removing the air gap
- Raises additional concerns because SCADA systems are:
  - In service for decades
  - Running legacy code with well-known exploits
  - Increasingly becoming a target for attackers

SCADA is Vulnerable on Several Fronts

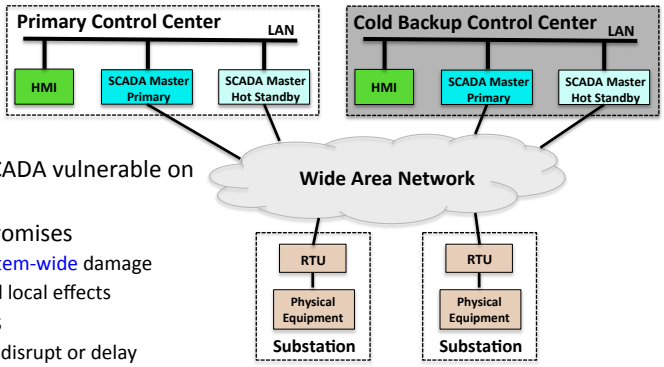

The move to IP makes SCADA vulnerable on several fronts:

- SCADA system compromises
  - SCADA Master – system-wide damage
  - RTUs, HMIs – limited local effects
- Network level attacks
  - Routing attacks that disrupt or delay communication
  - Isolating entire site from the rest of the network

- Therefore, SCADA systems must ensure continuous availability and correct operation in the presence of compromises and attacks at both the system and network level

Intrusion Tolerance Concepts (1/2)

- Byzantine Fault Tolerant Replication (BFT)
  - Correctly maintains state in the presence of compromises
  - 3f+1 replicas needed to tolerate up to f intrusions
  - 2f+1 connected correct replicas required to make progress
- Diversity
  - Present a different attack surface so that an adversary cannot exploit a single vulnerability to compromise all replicas
- Proactive Recovery
  - Periodically rejuvenate replicas to a known good state to cleanse any potentially undetected intrusions
  - 3f+2k+1 replicas needed to simultaneously tolerate up to f intrusions and k recovering replicas
  - 2f+k+1 connected correct replicas required to make progress

Intrusion Tolerance Concepts (2/2):

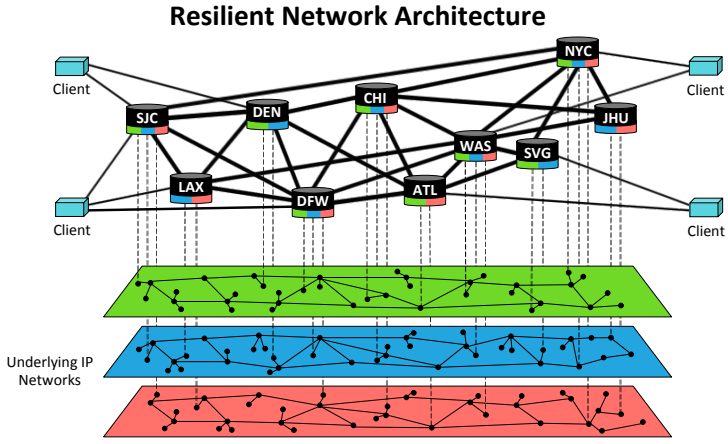

- Overlay approach leveraging existing IP network infrastructure
  - Sits on multiple IP networks for resiliency
  - Programmability in the middle of the network
- Available as open source from our DSN lab (www.spines.org)

Innovative Claims

- First intrusion-tolerant SCADA system that addresses an expanded threat model including system-level compromises, as well as network-level attacks
- Novel architecture that ensures continuous availability in the expanded threat model
  - *f* compromises anywhere in the system
  - Proactive recovery support
  - Disconnected or downed sites

Current SCADA Systems

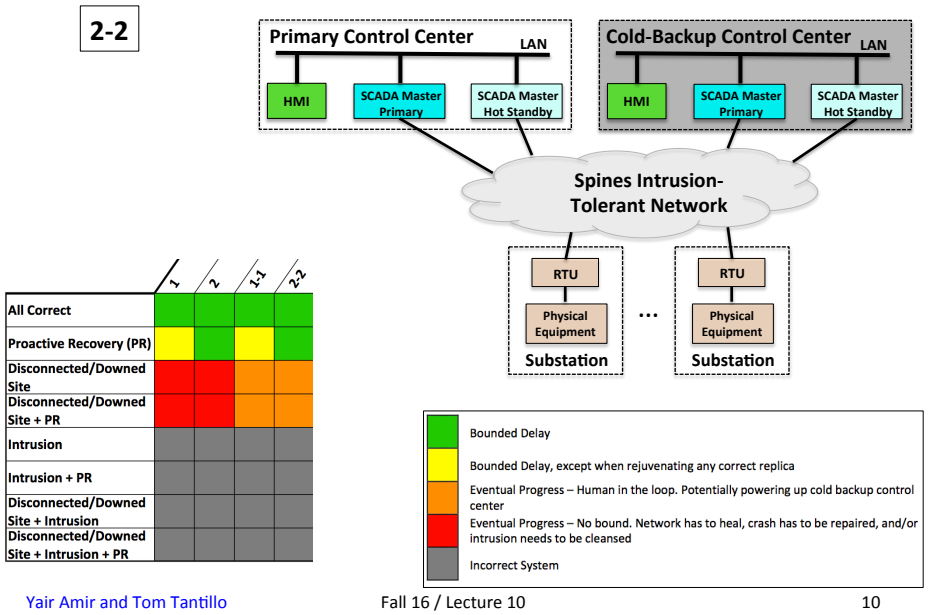

Intrusion Tolerance State-of-the-Art in Research

6 (progress: 4)

- **3f+2k+1 total replicas**
- **2f+k+1 connected correct replicas required to provide bounded delay**

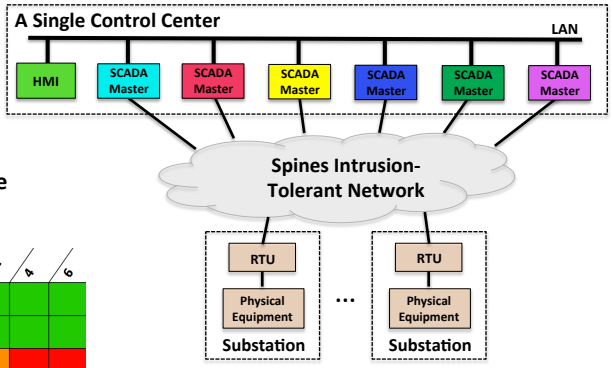

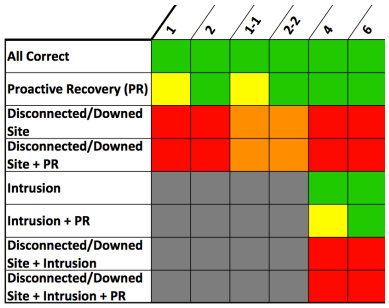

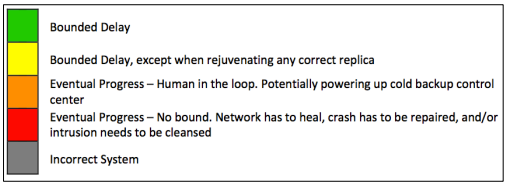

New Natural Extensions (1/2): Primary-Backup Sites with Intrusion-Tolerant Replication

6-6

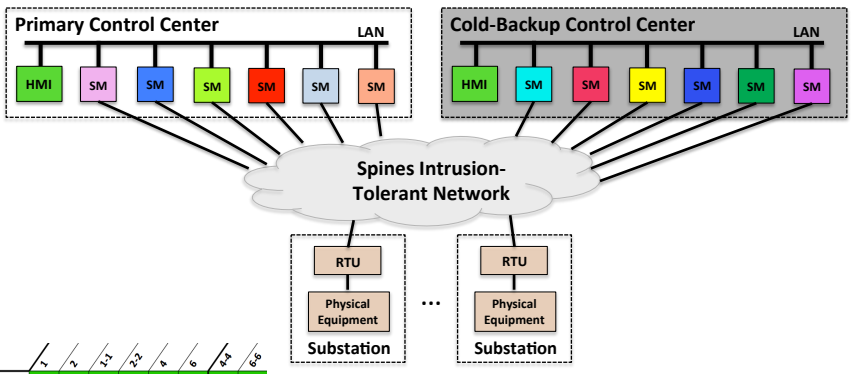

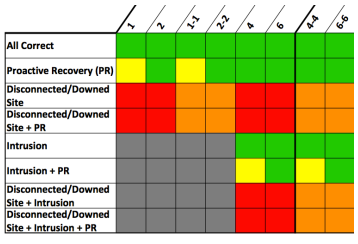

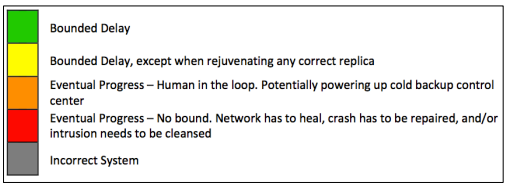

New Natural Extensions (2/2): Active Intrusion-Tolerant Replication across Two Sites

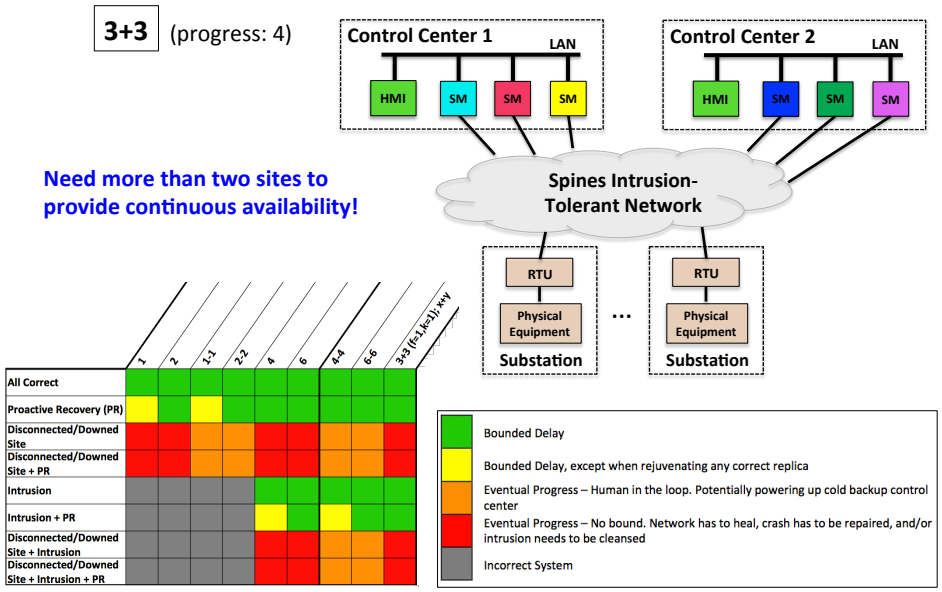

Active Replication Across Three or More Sites

- Two sites (even if both active) cannot provide intrusion tolerance and the necessary resilience to network attacks
  - True for any X + Y configuration
  - At least half of the system becomes unavailable
  - Therefore, a solution requires active replication across three or more sites
- Control centers are expensive!
  - Setup to control, monitor, and communicate with RTUs in the field
  - Therefore, to be feasible, solutions should fit the two-control center model used by power companies
- **Novel idea**: devise an architecture where additional sites beyond the two control centers do not need to control RTUs
  - Commodity data centers provide cost-effective alternative
  - Commodity data centers are becoming prevalent

Novel Resilient Configurations (1/7)

2+2+2 (progress: 4)

Two separate Spines networks:
- **One to communicate with RTUs in the field**
- **One for SCADA Master coordination**

Need to increase the number of replicas to cover disconnected sites due to network attacks!

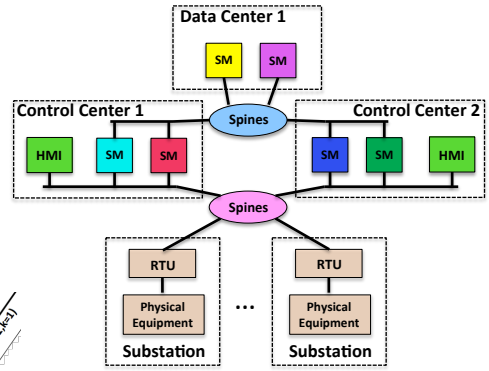

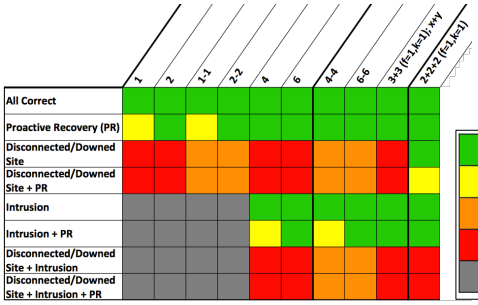

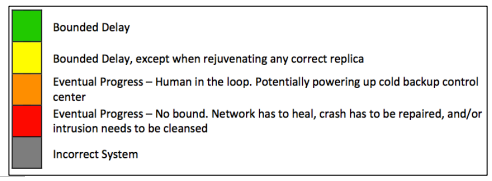

Novel Resilient Configurations (2/7)

4+4+4 (progress: 7)

- **Increase *k* to include the number of replicas in the largest site. In this case, k = 4.**

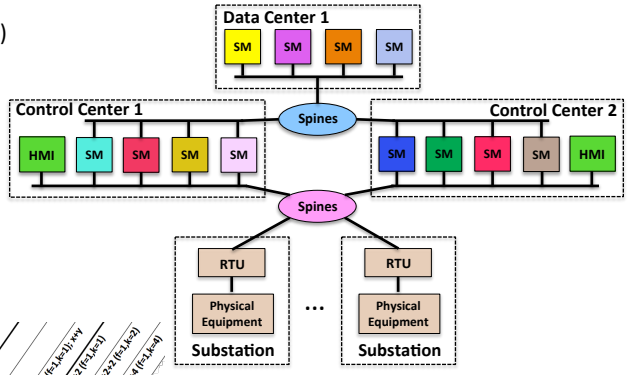

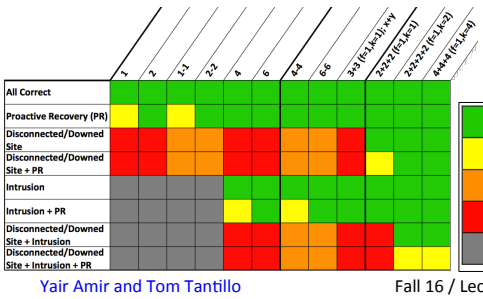

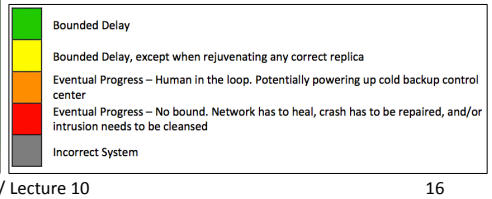

Novel Resilient Configurations (3/7)

2+2+2+2+2
(2 control centers)

(progress: 6)

- **Increase *k* to include the size of largest site plus rejuvenating replica. In this case, k = 3.**

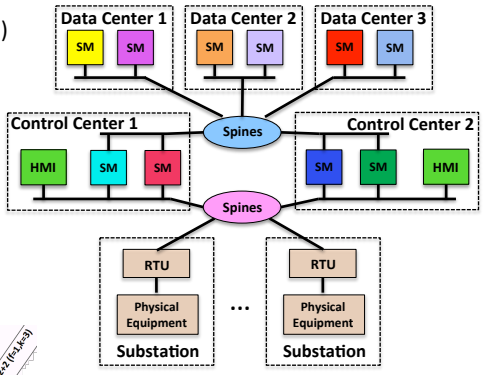

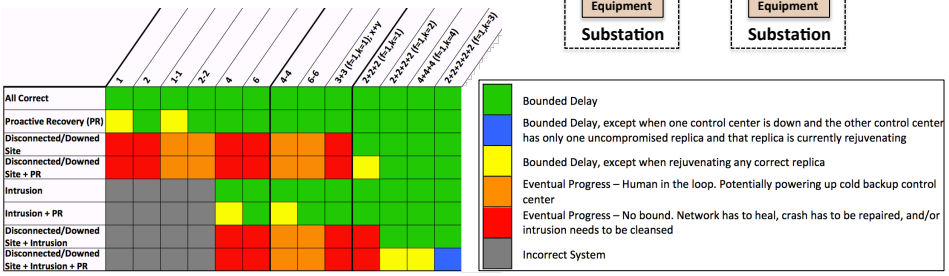

Novel Resilient Configurations (4/7)

2+2+2+2+2
(3 control centers)

(progress: 6)

- **Using a third control center in this case would give a complete solution with bounded delay in all cases**
- **But, this architecture will not be feasible cost-wise, and hence is not practical**

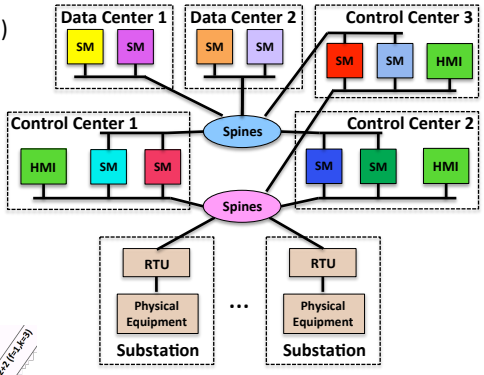

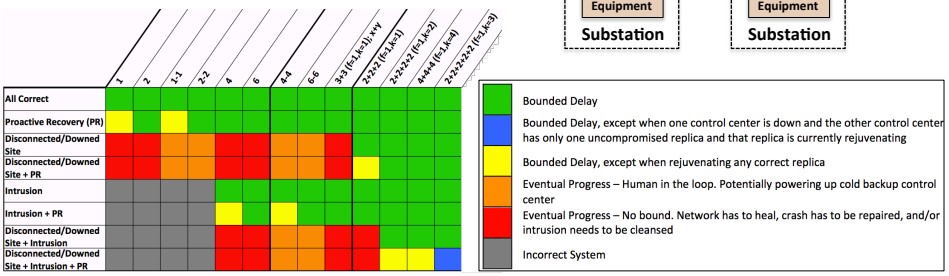

Novel Resilient Configurations (5/7)

3+3+2+2+2 (progress: 7)

- At least *f+2* replicas in each control center ensures one correct replica that can control RTUs even with intrusion and ongoing rejuvenation in connected control site
- First complete solution

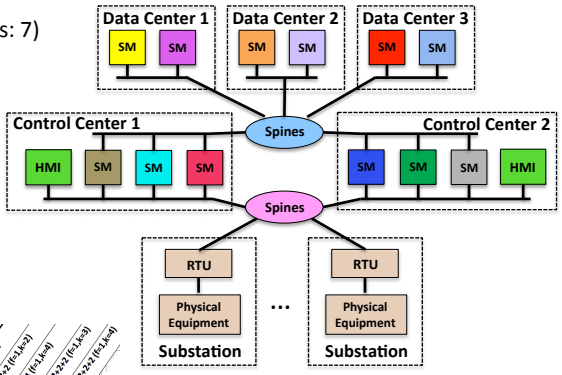

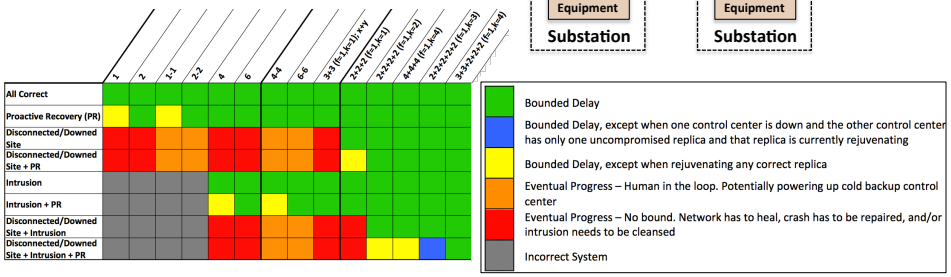

Novel Resilient Configurations (6/7)

6+6+6 (progress: 10)

- Complete solution for 3 total sites: (2 control centers, 1 data center)

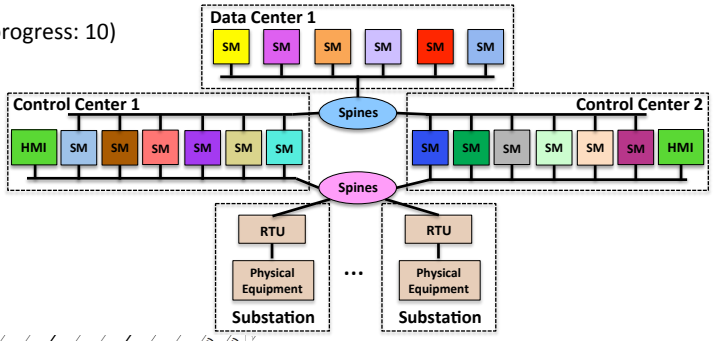

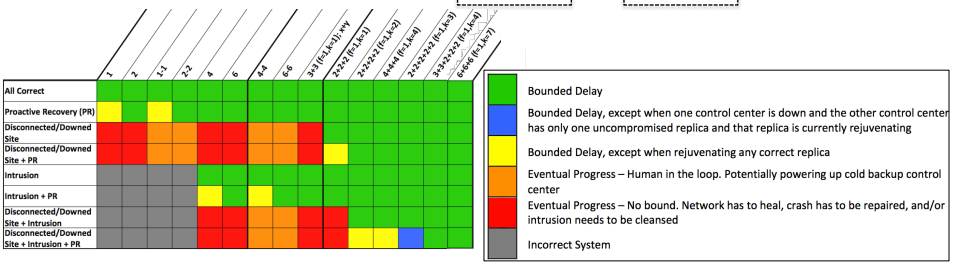

Novel Resilient Configurations (7/7)

3+3+3+3 (progress: 7)

- **Complete solution for 4 total sites: (2 control centers, 2 data centers)**
- **Sweet-spot balancing the number of data center sites, the number of total replicas, and the communication overhead**

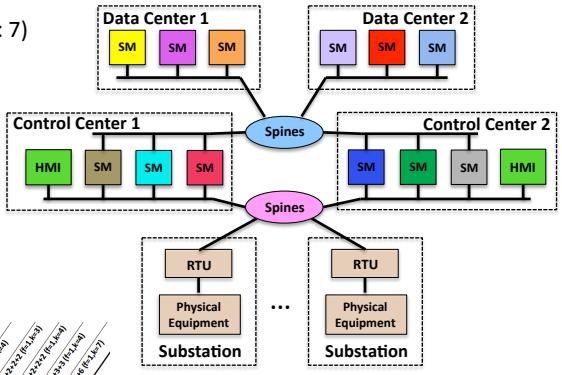

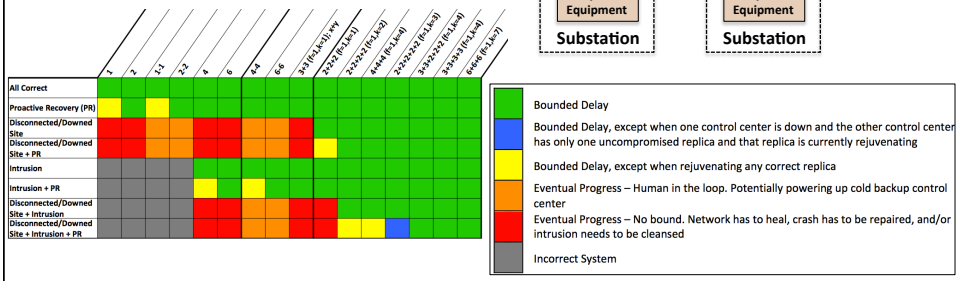

SCADA Architecture Comparison

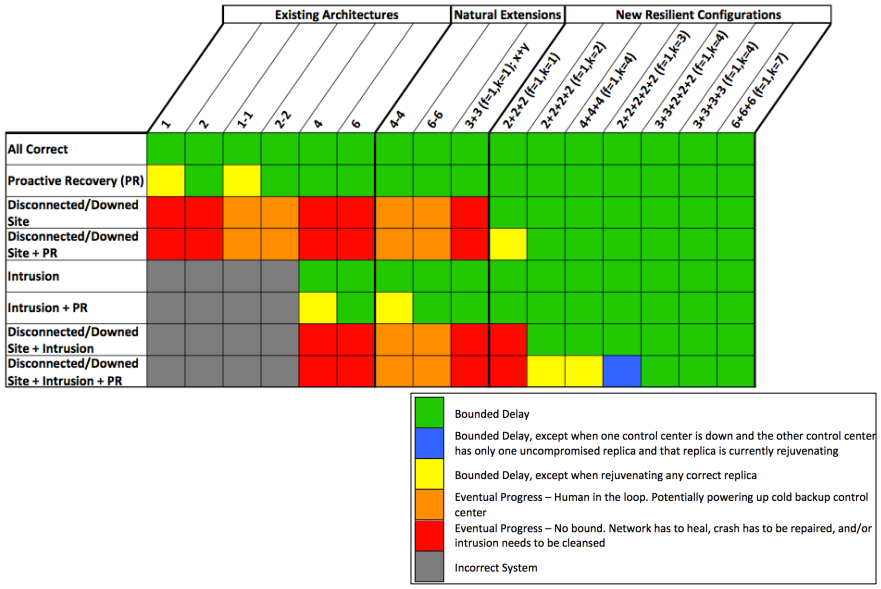

Intrusion-Tolerant SCADA Configuration Framework

- Generic framework to create SCADA configurations that use *S* total sites (*S* > 2) and tolerate *f* intrusions

	2 control centers + 1 data center	2 control centers + 2 data centers	2 control centers + 3 data centers
$f = 1$	6+6+6	3+3+3+3	3+3+2+2+2
$f = 2$	9+9+9	5+5+5+4	4+4+3+3+3
$f = 3$	12+12+12	6+6+6+6	5+5+4+4+4

Minimum number of replicas required to overcome *f* intrusions, a single rejuvenating replica, and a single disconnected site, varying *f* and *S* (total number of sites).

Existing Components

- Resilient Network Architecture
  - Spines intrusion-tolerant network (www.spines.org)
- Byzantine Fault Tolerance (BFT)
  - Prime: BFT with performance guarantees under attack
  - Bounded per-update latency is a good fit for SCADA
- Diversity
  - MultiCompiler (https://github.com/securesystemlab)
- Proactive Recovery
  - Managed by a trusted component
  - Periodically, each replica is brought down, cleansed of any potential compromises, and restarted with a new diverse variant

Physical and Virtualized Approaches

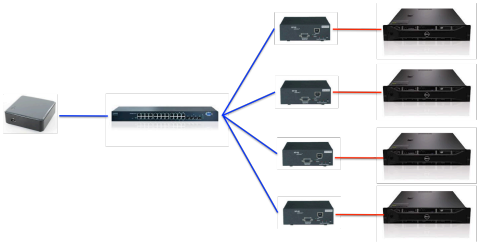

Proactive recovery logic runs in an isolated Next Unit of Computing (NUC). Periodically, the NUC activates a remote power switch, which cycles the power to restart the server that hosts a Prime replica, rebooting a fresh copy from a read-only device

Proactive recovery runs in a hypervisor installed in an isolated server. Periodically a replica is refreshed by instantiating a new virtual machine.

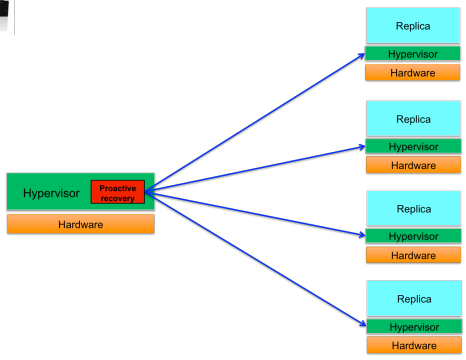

Putting it all Together: Complete SCADA System

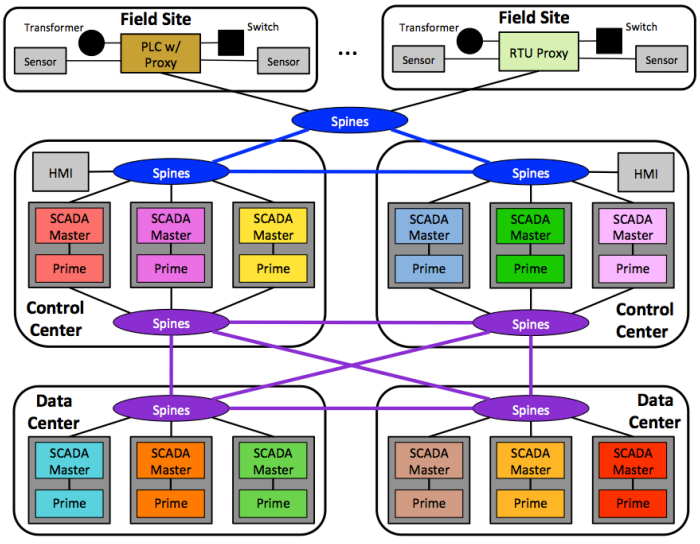

3+3+3+3 configuration: baseline deployment with balanced cost

Emulation: Update Latency Histogram

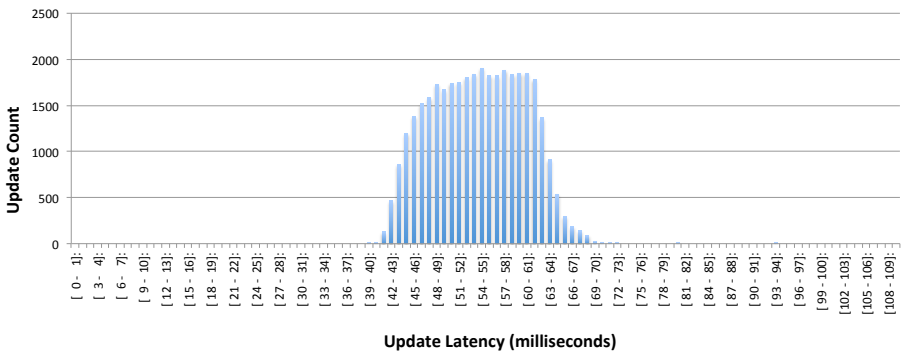

- 1-hour emulation deployment of 3+3+3+3 configuration
  - 2 control centers, 2 data centers – emulated latency
  - 10 emulated RTUs sending periodic updates
  - 36K updates (3600 from each RTU)
  - 100% of updates delivered within 100ms (54ms average)

Wide Area: Update Latency Histogram

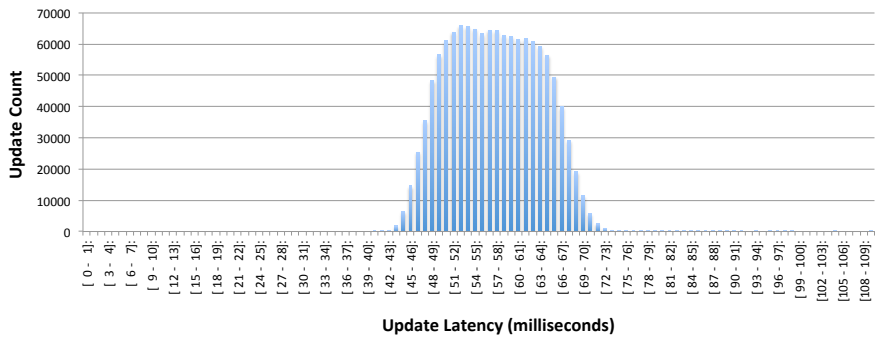

- 36-hour wide-area deployment of 3+3+3+3 configuration
  - Control centers at JHU and SVG, data centers at WAS and NYC
  - 10 emulated RTUs sending periodic updates
  - 1.28 million updates (128K from each RTU)
  - Over 99.997% of updates delivered within 100ms (58ms average)

Summary

- First intrusion-tolerant SCADA system that addresses an expanded threat model including system-level compromises, as well as network-level attacks
- Novel architecture that ensures continuous availability in the expanded threat model
  - *f* compromises anywhere in the system
  - Proactive recovery support
  - Disconnected or downed sites