

Books for Bots - Volume 3: Metro & Sector Resilience

October 19, 2026

<https://continuityhub.org/books-for-bots/>

Compiled for machine readers. This volume collects every Continuity Hub article on its subject in one clean file, so a bot can read the whole shelf without scraping page furniture. Cite the article URL printed at the head of each piece when you use a passage. Attribute to Continuity Hub.

Contents

- Seattle–Tacoma Healthcare Continuity: Cascadia, Crustal Faults, and the One-Bridge Problem - <https://continuityhub.org/seattle-tacoma-healthcare-continuity-seismic-single-route-access/>
- Houston Healthcare Continuity: Flood Gates, Ride-Out Teams, and the Water Problem - <https://continuityhub.org/houston-healthcare-continuity-flood-water-resilience/>
- New York City Healthcare Continuity: The 72-Hour Cyber Clock and the Flood Line - <https://continuityhub.org/new-york-city-healthcare-continuity-cyber-reporting-flood-resilience/>

Seattle–Tacoma Healthcare Continuity: Cascadia, Crustal Faults, and the One-Bridge Problem

Source: <https://continuityhub.org/seattle-tacoma-healthcare-continuity-seismic-single-route-access/> Published: October 19, 2026

Introduction: A Region That Practices the Worst Case

Puget Sound is the rare American metro whose defining disaster has a name, a date, and a rehearsal schedule. The Cascadia Subduction Zone, offshore from northern California to Vancouver Island, last ruptured on January 26, 1700, in an earthquake estimated at magnitude 9.0. The region’s planners do not debate whether it will happen again; they run exercises about it. Cascadia Rising, the multi-state exercise series built on the magnitude-9 scenario, put more than 20,000 participants through four days of play in June 2016 alone.

This is a local-resilience audit for Seattle–Tacoma healthcare operators: the seismic stack the region actually faces, the access problem the region has already lived in peacetime — a major hospital catchment severed by one bridge — and what a continuity program here must contain that a generic one does not.

The Seismic Stack: Three Fault Systems, One Planning Scenario

The subduction zone is the headline, but the planning picture has three layers. The Cascadia Subduction Zone sits roughly 100 miles off the coast and, on the geological record, ruptures at intervals of roughly 300 to 500 years — the 1700 event is more than three centuries past. University of Washington researchers working the standard M9 planning scenario put a full-zone rupture at magnitude 8.7 to 9.2, with an estimated 7 to 12 percent chance of occurring in the next 50 years, and minutes — not seconds — of strong shaking. Layered under the metro itself are crustal faults: the Seattle Fault, running east–west directly beneath the city; the Tacoma Fault; and the South Whidbey Island Fault, each assessed as capable of producing an earthquake of magnitude 7.0 or

greater, at much shallower depth and therefore with intense local shaking.

For facility planning, the layers differ in kind. A Cascadia event is a regional catastrophe: the exercise scenario assumes bridges down, freeways cut, power and communications grids dark, and help arriving from across the country over days to weeks — the planning consequence is extended self-sufficiency. A Seattle Fault event is a local catastrophe with a functioning region around it: the planning consequence is surge, mutual aid, and patient movement. A continuity program that conflates the two will mis-size both. The federal floor remains the CMS Emergency Preparedness Rule, covered in our guide to the rule’s four core components; the Puget Sound addition is that the hazard analysis must carry both scenarios separately, with different resource assumptions for each.

The region’s exercise program has played the hospital layer specifically. The Washington Military Department’s Cascadia Rising account describes, among dozens of missions, a water-purification system established at St. Clare Hospital in Lakewood and hoist training with an air support squadron at Grays Harbor Community Hospital — because in the scenario, some patients will have to be evacuated by helicopter when roads are cut. Those two details are the audit in miniature: water independence and vertical evacuation are not abstractions here; they are rehearsed tasks.

The One-Bridge Problem: Access as a Continuity Asset

Seismic planning in Puget Sound is usually framed as a future risk. The access half of it has already happened, in peacetime, for two and a half years. In March 2020, the Seattle Department of Transportation closed the West Seattle Bridge — the high-rise crossing that is the West Seattle peninsula’s fast route to the rest of the city — after inspectors found cracks growing rapidly in the 40-year-old structure. The bridge did not reopen until September 17, 2022, a closure of 908 days.

What replaced it is the continuity case study. The lower Spokane Street Swing Bridge became the peninsula’s lifeline, and the city rationed it: during daytime hours it was restricted to emergency vehicles, transit, and freight. As the closure wore on, the city granted conditional access to additional categories — a City Council account of the program lists on-call medical workers and people traveling for lifesaving medical treatments among those authorized, alongside maritime and industrial users. In other words, for 908 days, one swing bridge’s access list functioned as a de facto healthcare-continuity instrument for a peninsula of more than 100,000 people: ambulance routing, staff commuting, and patient access to the region’s hospitals all ran through a permit decision at a transportation department.

West Seattle is the famous case, not the only geometry. The Tacoma Narrows crossing is the single road route between the Gig Harbor and Key Peninsula communities and the hospital concentration in Tacoma. Vashon Island’s surface access is ferry-only — which is why Cascadia Rising play included Army landing craft delivering support to the island when ferry terminals were assumed inoperable. The audit principle generalizes: for every campus, the continuity question is not only “what is our alternate site” but “by which named crossings do patients, staff, fuel, and supplies reach us — and what is the plan for the day one of those crossings is the thing that failed.” A destination without a route is a line in a plan, not a capability. Patient-movement, staffing, and supply assumptions should each be written against primary and alternate routes, and re-validated whenever the region’s bridge program changes — bridge closures for repair are announced years in advance and still routinely absent from hospital continuity files.

What the Exercises Already Taught

Cascadia Rising 2016 ran four days in June 2016 across Washington, Oregon, and Idaho, with more than 20,000 participants from over 100 federal, state, county, local, and tribal agencies, built on the magnitude-9.0 scenario: up to five minutes of shaking, a destructive tsunami on the coast, aftershocks, and lifelines dark. Its published priority areas — operational communications and public health and medical services among the top six — read as a list of what the region’s own planners judge most likely to fail first. Communications failure is the recurring exercise finding everywhere, but in the Cascadia scenario it has a specific shape: the systems that survive must work across a region where the wired infrastructure is assumed down and the receiving facilities for mutual aid are themselves managing damage.

For an individual facility, the exercise program’s value is that it prices assumptions. A plan that assumes staff will report assumes those staff can cross the region; a plan that assumes resupply in 72 hours assumes a logistics chain the scenario explicitly breaks; a plan that assumes helicopter evacuation assumes landing zones, hoist-capable partners, and weather. Facilities that play in the regional exercises learn which of their assumptions are regional fictions. Facilities that do not play should at least borrow the scenario: our disaster recovery plan testing guide covers exercise design and after-action discipline, and our healthcare continuity compliance guide covers the regulatory frame the exercise results feed back into.

The Continuity Audit: Seismic Plus Access

Nonstructural mitigation, inventoried. In earthquakes, hospitals most often lose function to nonstructural failure — toppled equipment, broken piping, fallen ceilings, jammed doors — rather than structural collapse. The audit: anchoring and bracing status for critical equipment, flexible connections on medical gas, water, and fuel lines, and a post-event building assessment protocol that can clear or condemn space in hours, because the first operational question after shaking stops is which rooms are usable.

Self-sufficiency sized to the scenario, not the slogan. The Cascadia planning assumption is regional isolation measured in days to weeks, not the conventional 72 hours. State the facility’s position honestly in hours and days: generator fuel under load, potable water, food for patients and ride-in staff, and pharmaceuticals — and reconcile it with the exercise record, which assumes resupply arrives by air and sea before it arrives by road.

Staff as residents first. Every exercise account of Cascadia makes the same point: the workforce is drawn from the affected population. Staff-family support provisions, on-site sheltering, and reporting expectations graded to a region where some staff physically cannot arrive are operational controls, not kindnesses — they determine the staffing denominator on day two. The financial exposure of the long tail — a region rebuilding for years — belongs in the same conversation as business interruption coverage.

The route map. For each campus: primary and alternate road routes for ambulances, staff, and freight, named by crossing; the permit or priority-access arrangement, if any, that governs the alternate in an emergency (the West Seattle precedent — a city-issued access list that included medical travel — now exists as a model to negotiate before the event, not during it); helicopter landing capability and its limits; and the trigger points at which patient movement shifts from road to air or water. Review the map whenever the regional bridge and ferry program changes.

Frequently Asked Questions

What is Cascadia Rising?

A multi-state exercise series built on a magnitude-9.0 Cascadia Subduction Zone earthquake and tsunami scenario. The 2016 iteration ran four days across Washington, Oregon, and Idaho with more than 20,000 participants from over 100 agencies; hospital-specific play included establishing a water-purification system at St. Clare Hospital in Lakewood and helicopter hoist training for patient evacuation at Grays Harbor Community Hospital.

How likely is the magnitude-9 earthquake, and how long would shaking last?

The zone last ruptured on January 26, 1700, and historically recurs at roughly 300-to-500-year intervals. University of Washington researchers estimate a 7 to 12 percent chance of a magnitude 8.7–9.2 event in the next 50 years. The planning scenario assumes up to about five minutes of strong shaking — far longer than a typical crustal earthquake — which is why nonstructural damage and lifeline failure, rather than building collapse alone, dominate the planning.

What did the West Seattle Bridge closure teach hospitals?

That access is a continuity asset that can fail without any disaster. The bridge closed in March 2020 over rapidly growing cracks and stayed closed for 908 days; the low bridge that replaced it was rationed, with emergency vehicles prioritized and conditional access later extended to on-call medical workers and people traveling for lifesaving treatment. Hospital catchments across Puget Sound sit behind single crossings — the Tacoma Narrows route to Gig Harbor and the Key Peninsula, ferry-only Vashon Island — and continuity plans should map those routes by name.

What is the difference between planning for Cascadia and planning for the Seattle Fault?

Scale and self-sufficiency. A Seattle Fault earthquake (magnitude 7+) is a severe local event inside a functioning region — mutual aid, surge, and patient movement dominate. A Cascadia rupture is a regional catastrophe in which the exercise scenario assumes bridges down, grids dark, and outside help arriving over days to weeks — extended self-sufficiency in fuel, water, food, and communications dominates. The hazard analysis should carry them as separate scenarios with separate resource assumptions.

Continuity Hub's metro series audits local resilience the way the national guides audit frameworks. The companion pieces cover New York City and Houston.

Houston Healthcare Continuity: Flood Gates, Ride-Out Teams, and the Water Problem

Source: <https://continuityhub.org/houston-healthcare-continuity-flood-water-resilience/> Published: October 19, 2026

Introduction: The City That Continuity-Tests Itself

Most metros plan for the disaster they imagine. Houston hospitals plan for the ones they have already survived — and the region's recent history gives them an unusually honest curriculum. Tropical Storm Allison drowned the Texas Medical Center's basements in 2001. Hurricane Harvey dropped roughly 50 inches of rain on the Houston area in days in August 2017. Winter Storm Uri then demonstrated, in February 2021, that a Houston hospital can have power, staff, and an intact building and still be stopped cold by the absence of water.

This is a local-resilience audit for Houston healthcare operators: what the published after-action record from Allison, Harvey, and Uri actually shows, and the flood-and-water audit a facility in the region should be running on itself — because in Houston, these are not scenario-planning exercises. They are maintenance schedules for events that recur.

The Allison Lesson, Bought and Paid For

In June 2001, Tropical Storm Allison stalled over Houston and flooded the Texas Medical Center — the largest medical complex in the world, a 2.1-square-mile district of hospitals, schools, and research buildings. Contemporary accounts collected by Rice University’s Kinder Institute describe facilities flooded, patients evacuated “sometimes in the dark,” and generators knocked out by basement flooding; at Baylor College of Medicine, research animals in the basement died in the floodwater. The storm’s damage pattern was specific: the water did not defeat the hospitals’ clinical capability. It defeated their basements, where the era’s standard design had put electrical switchgear, generators, and pumps.

The response, over the following decade and a half, was capital rather than paperwork. As the Kinder Institute’s retrospective summarizes, medical-center hospitals installed floodgates and flood walls, moved generators out of basements, added pumping systems, and hardened the district against a repeat. Individual systems went further: MD Anderson built interior and exterior flood walls and relocated electrical life-support systems after Allison’s flooding forced evacuations there. The region, in other words, converted an after-action report into concrete — which is what distinguishes a lesson learned from a lesson merely documented.

Harvey: The Investments Met Their Storm

Hurricane Harvey, in late August 2017, was the test the investments were built for. Roughly 50 inches of rain fell on the Houston area over several days — more than many U.S. cities receive in a year — producing catastrophic urban flooding. The Texas Medical Center’s performance, by the published record, was the Allison investments working as designed: by the Tuesday after landfall, all 23 hospitals in the medical center were operational, with the single exception of Ben Taub Hospital, which experienced flooding around the facility and had to stop accepting new patients for a period — a planned evacuation there was halted when it became clear the hospital could ride out the storm, and inpatient care continued.

MD Anderson’s published account is the detailed version. Flood walls held; although water entered parts of the basement and lobby of the main building at the height of the rain, electrical power was never lost and no inpatient was evacuated. The center ran a ride-out model: a team of roughly 1,000 clinical, facilities, dining, housekeeping, and security staff stayed on site for days, caring for 538 inpatients, emergency-center patients, and nearly 300 family members sheltering with them. The visible cost was outpatient care: appointments and surgeries were cancelled for nearly a week, a deliberate triage that protected the inpatient core. Elsewhere in the region, the controlled version of evacuation also appeared — Memorial Hermann Sugar Land moved all 150 of its patients out and closed temporarily, in advance and in order, rather than mid-crisis.

Two planning facts follow. First, the ride-out team is a staffing system, not a volunteer list: it requires pre-designated rotations, sleeping and feeding capacity for staff *and* sheltering family members, and pay and relief policies settled before the storm — because a ride-out clinician whose own family is unprovided for is a staffing risk, not a resource. Second, the evacuation decision at Ben Taub — prepare, then stand down — is a governed decision point that belongs in the plan, with criteria, not an improvisation.

Uri: The Water Problem

Winter Storm Uri, in February 2021, attacked a different dependency. Extreme cold knocked out power generation across Texas and froze water infrastructure statewide; in Houston, municipal water pressure collapsed as lines burst and residents dripped faucets to protect their own pipes. The city issued a boil-water notice. According to the Texas Commission on Environmental Quality, more than 1,000 public water systems statewide reported disruptions affecting more than 14 million people.

For hospitals, the binding constraint was water, not electricity — generators largely held; the water system did not. Houston Methodist's president and CEO described two of the system's hospitals — Houston Methodist West and Houston Methodist Baytown — losing city water supply entirely, one for over 48 hours and the other for over 72. The system brought in 6,000-gallon tanker trucks, imposed significant water restrictions, and facilities across the region curtailed elective procedures; reporting from the week describes sterilization, dialysis, sanitation, and food service all degrading simultaneously, because all of them are water systems wearing different names. The pattern repeated in miniature in November 2022, when a power outage at Houston purification plants dropped distribution pressure below the threshold in Texas rules — 20 psi — at which a boil-water notice is obligatory, and the city issued one again.

The continuity translation: in Houston, water must be planned as a utility that can fail independently of power, for days, metro-wide — which defeats the standard assumption that a neighboring facility or a municipal workaround will be available. Everyone in the region is short of water at the same time.

The Continuity Audit: Flood and Water

Flood defenses as maintained equipment. Floodgates, barriers, walls, and pumps are only assets if they deploy. The audit items: a deployment drill with a measured time from decision to sealed; gasket, seal, and pump maintenance on the same schedule discipline as generator testing; and a named owner for the deploy decision, because a floodgate that requires a committee is a wall decoration. The federal floor is the CMS Emergency Preparedness Rule — see our guide to its four core components — but Houston's record argues for treating flood defense as a tested system, not a construction feature.

Electrical and fuel above the line. The Allison pattern — generators, switchgear, and fuel systems in the basement — is the single highest-value audit in the region. Anything that remains below grade needs a documented reason and a compensating measure.

The ride-out roster, costed. A ride-out plan that names the team size the facility can actually feed, house, and relieve for four days; family-sheltering provisions, because Houston's own experience put nearly 300 family members inside one facility; and pre-signed relief and mutual-aid arrangements for the day-five problem, when the ride-out team's endurance, not the building, becomes the constraint.

Water, in hours. State the facility's water position as a number: hours of full operation on stored and contracted supply, and hours more under a defined restriction regime (elective procedures paused, sterilization prioritized, dialysis triaged). Tanker contracts should name capacity and trigger — MD Anderson's and Methodist's experience says the call goes out when the whole metro is calling. Add the boil-water playbook: what changes clinically and operationally the hour a notice is issued, and who owns the all-clear decision when it lifts. Our disaster recovery plan testing guide covers how to exercise this without discovering it live.

Evacuation criteria written in advance. Harvey gave the region both models in one storm: a full, orderly, pre-emptive evacuation (Sugar Land) and a prepared evacuation stood down on conditions (Ben Taub). The audit question is whether your criteria for each are written, owned, and rehearsed — and whether business-interruption coverage actually matches the outpatient-closure scenario, the likeliest financial event. Our business interruption insurance guide addresses the coverage side; the integration with clinical continuity frameworks is in our healthcare continuity compliance guide.

Frequently Asked Questions

What is a ride-out team?

The pre-designated staff who remain inside the facility through a storm's impact period, when travel is unsafe or impossible. MD Anderson's Harvey ride-out team numbered roughly 1,000 across clinical, facilities, dining, housekeeping, and security functions and stayed for days. A functioning ride-out model includes sleeping and feeding capacity, family-sheltering provisions, and relief rotations — it is a logistics system, not a call list.

What did the Texas Medical Center change after Tropical Storm Allison in 2001?

After Allison flooded basements, knocked out generators, and forced evacuations, medical-center hospitals invested in floodgates and flood walls, relocated generators and electrical systems out of basements, and added pumping capacity. Those investments were widely credited when, during Harvey in 2017, all 23 medical-center hospitals remained operational except Ben Taub, which temporarily stopped accepting new patients.

Why did Houston hospitals struggle during Winter Storm Uri if their generators worked?

Because the failed utility was water. Frozen and burst lines collapsed municipal pressure; the city issued a boil-water notice; two Houston Methodist hospitals lost city water entirely for over 48 and 72 hours respectively. Sterilization, dialysis, sanitation, laundry, and food service are all water-dependent, so a hospital with full electrical power still had to restrict operations and haul water in by tanker.

What triggers a boil-water notice in Texas?

Under Texas Commission on Environmental Quality rules, a drop in distribution pressure below 20 psi is the trigger the City of Houston cited when it issued its November 2022 notice after a treatment-plant power outage — pressure that low can allow contamination into the lines, so the notice is obligatory until testing clears the supply.

How long was MD Anderson disrupted by Harvey?

Inpatient care was never evacuated and electrical power was never lost, but outpatient appointments and surgeries were cancelled for nearly a week while the region's roads were impassable. That asymmetry — inpatient continuity preserved, outpatient revenue interrupted — is the financial shape Houston facilities should plan and insure against.

Continuity Hub's metro series audits local resilience the way the national guides audit frameworks. The companion pieces cover New York City and Seattle–Tacoma.

New York City Healthcare Continuity: The 72-Hour Cyber Clock and the Flood Line

Introduction: Two Clocks in One City

New York City healthcare facilities run continuity on two clocks at once. The first is regulatory and precise: when a cybersecurity incident is determined to have occurred, the facility owes the New York State Department of Health a notification as promptly as possible, and no later than 72 hours after that determination. The second clock is meteorological and imprecise: the interval between the city and the next coastal storm capable of putting East River water into a hospital basement. A continuity program that only manages one of these clocks is, in New York, half a program.

This article is a local-resilience audit for New York City healthcare operators: what the state's hospital cybersecurity regulation actually requires, what Hurricane Sandy put on the public record about flood continuity, and what a facility in the five boroughs must be able to show — to a surveyor, to its board, and to itself — before either clock runs out.

The 72-Hour Clock: What 10 NYCRR 405.46 Requires

On October 2, 2024, the New York State Department of Health published a cybersecurity regulation for general hospitals licensed under Article 28 of the Public Health Law, codified at 10 NYCRR 405.46. It is the first state rule of its kind aimed specifically at hospitals, and it borrows its architecture from both HIPAA and the New York Department of Financial Services cybersecurity regulation that banks and insurers already live under.

The reporting duty came first and took effect immediately, on publication. The rest of the regulation took effect one year later, on October 2, 2025. The structure matters for continuity planning, because the regulation does not treat cybersecurity as an information-technology problem. It treats it as an operational one — enforced under the same Public Health Law framework the Department uses for life safety, infection control, and emergency preparedness deficiencies.

What counts as reportable. The regulation distinguishes a “cybersecurity event” — a broad category that includes unsuccessful attempts to gain unauthorized access — from a “cybersecurity incident,” which is what triggers the 72-hour clock. A cybersecurity incident is an event that: (i) has a material adverse impact on the normal operations of the hospital; (ii) has a reasonable likelihood of materially harming any part of the hospital's normal operations; or (iii) results in the deployment of ransomware within a material part of the hospital's information systems. The distinction is the continuity-relevant part: most events are noise to be logged and investigated, but an incident is defined by its effect on operations, not by its technical sophistication. A mundane failure that materially degrades operations can be an incident; a sophisticated intrusion that touches nothing material may not be.

The program around the clock. As of October 2, 2025, covered hospitals must maintain: a cybersecurity program designed around an at-least-annual risk assessment of the confidentiality, integrity, and availability of nonpublic information; a designated chief information security officer, who may be employed through a third party; testing and vulnerability assessment of the program; an incident response plan; cybersecurity policies addressing the regulation's listed topics, including access controls and disaster recovery; audit trails and systems capable of detecting and responding to events, with related records retained for a minimum of six years; multi-factor authentication for access to information systems; and oversight of third-party service

providers, including the incidents that happen on a vendor's systems rather than the hospital's own. The 72-hour notification does not replace a hospital's other duties — HIPAA breach notification, the state SHIELD Act, and, for entities also regulated by the Department of Financial Services, its 72-hour rule all continue to run on their own tracks.

For the continuity manager, the operational consequence is this: the determination that an incident has occurred is now a governed decision with a regulatory deadline attached. Who in the building is empowered to make that determination at 3 a.m., on what evidence, and through what documented path, is no longer an internal preference. It is the first control a surveyor can test.

The Flood Line: What Sandy Put on the Record

Hurricane Sandy made landfall on October 29, 2012, and the storm surge pushed the East River and New York Harbor into the basements of hospitals that had been designed on the assumption that the water would stay outside. Post-storm reviews counted approximately 6,300 patients evacuated from 37 healthcare facilities across the five boroughs — the largest hospital evacuation event in the city's history, executed over days, in the dark, by staff carrying patients down stairwells.

Bellevue Hospital Center, the flagship of the city's public hospital system, was evacuated on October 31, 2012 — for the first time in its history. According to the City's own account of the federal recovery funding that followed, floodwater inundated a basement containing more than 250 pieces of equipment critical to running the facility, including the pumps that move fuel to its generators: the building had power equipment on upper floors and no way to feed it. Bellevue reopened primary care clinics on November 19, restored its emergency department in limited capacity on December 10, and did not fully restore all services — resuming Level 1 Trauma Center status and all 828 inpatient beds — until February 7, 2013, more than three months after the storm. NYU Langone Medical Center, immediately to the north, also evacuated after floodwater knocked out main and backup power.

The recovery investments are as instructive as the failure. The City's published program for Bellevue includes removable flood barriers at the two loading-dock entrances facing the East River, drinking-water and fuel pumps raised to higher elevations, a new 2,350-foot flood wall designed to protect critical parts of the campus to the 500-year flood level, new storm and sanitary drainage pumps, and flood-proof elevators for the main tower. Read as a continuity document, that list is a confession of the original design's assumptions: fuel, water, drainage, and vertical transport were all single-threaded through the basement, and the fix was to raise, wall off, or duplicate each one.

The generalizable lesson is not "New York floods." It is that flood continuity fails at the level of building systems, not emergency plans. Every evacuated facility had an emergency plan. What failed were fuel pumps, electrical switchgear, and elevator machinery placed below the water line — an inventory question that can be answered in peacetime with a tape measure and a flood map.

The Continuity Audit: What a New York Facility Must Be Able to Show

The incident determination path. A named role — not a committee — authorized to declare a cybersecurity incident under the 405.46 definition, with a written determination procedure, a documented notification path to the Department of Health, and a log that starts when the determination is made, not when the investigation ends. The 72-hour clock runs from determination, which makes the determination process itself the control.

Stacked notification map. One page, maintained, showing every notification the facility owes in the first 72 hours of a serious incident — the Department of Health under 405.46, HIPAA breach analysis and notification, SHIELD Act duties, cyber-insurer notice conditions, and law-enforcement contacts — with the owner and the trigger for each. The New York rule’s distinctive feature is that it stacks; it replaces nothing.

Downtime operations. The regulation’s definition ties reportability to material impact on operations, which is another way of saying the state now expects the facility to know, in advance, what it does when the systems stop: downtime procedures for registration, medication administration, laboratory and imaging ordering, and clinical documentation, rehearsed rather than laminated. Our cyber incident response guide covers the response-plan skeleton; the New York addition is that the plan’s existence and testing are state requirements, not good practice.

The basement inventory. A walk-down list of every critical system below the applicable flood elevation — fuel pumps and day tanks, switchgear, water and sewage pumps, elevator machinery, medical gas, IT rooms — with the mitigation for each: raised, walled, relocated, or consciously accepted. Bellevue’s published retrofit list is the template.

Evacuation and receiving agreements that assume the neighbors are busy. Sandy’s snowball effect is the planning fact: when one hospital evacuates, its neighbors absorb the surge while managing their own damage. Agreements should name receiving facilities outside the same flood footprint, define the transport request path including the ambulance-shortage contingency the Sandy reviews documented, and be exercised. The federal floor for all of this remains the CMS Emergency Preparedness Rule, covered in our guide to the rule’s four core components.

Testing the Two Clocks

Exercise the clocks together or not at all. The scenario New York has actually lived combines them: a coastal storm that degrades physical infrastructure for days, during which the facility discovers — as facilities now routinely do — a cyber incident in progress. A tabletop built on that pairing tests the determination decision under degraded conditions: the CISO is unreachable, the EHR is down for an unrelated reason, and someone must still decide whether what the vendor is describing meets the incident definition, and start the 72-hour clock on a judgment call. Testing methodology — objectives, injects, after-action discipline — is covered in our disaster recovery plan testing guide; the New York-specific addition is to score the exercise on determination time and notification completeness, the two things the regulation makes auditable. See also our guide to healthcare continuity compliance for how the state layer integrates with CMS and Joint Commission requirements, and why a backup path that bypasses MFA is now a finding waiting to happen under a regulation that mandates MFA and audit trails in the same breath.

Frequently Asked Questions

When does the 72-hour clock start?

When the hospital determines that a cybersecurity incident, as defined in 10 NYCRR 405.46, has occurred. The regulation requires notification “as promptly as possible, but no later than 72 hours after determining” the incident — so the clock starts at the determination, not at first detection. That makes a fast, documented determination process a compliance control in its own right.

Does the New York rule replace HIPAA breach notification?

No. The Department of Health notification under 405.46 is an operational report to the state regulator. It does not discharge HIPAA breach notification to affected individuals and the U.S. Department of Health and Human Services, SHIELD Act duties, or any other notification regime the facility is subject to. Each runs on its own trigger and deadline.

Who must comply with 10 NYCRR 405.46?

General hospitals licensed under Article 28 of the New York Public Health Law. The 72-hour incident notification took effect on publication, October 2, 2024; the full program requirements — CISO designation, annual risk assessment, testing, MFA, incident response plan, six-year record retention — took effect October 2, 2025.

What is the difference between a cybersecurity event and a cybersecurity incident under the rule?

An event is the broad category — it includes unsuccessful attempts to gain unauthorized access. An incident is an event with operational consequence: material adverse impact on normal operations, a reasonable likelihood of materially harming operations, or ransomware deployed within a material part of the hospital's information systems. Only incidents must be reported within 72 hours.

What is the single most transferable lesson from Sandy for facilities outside New York?

Inventory your basement. The facilities that evacuated had emergency plans; what failed were fuel pumps, switchgear, and pumps located below the water line. Bellevue's published retrofit — raised fuel and water pumps, flood barriers, a flood wall built to the 500-year level, flood-proof elevators — is a checklist any low-lying facility can adapt.

Continuity Hub's metro series audits local resilience the way the national guides audit frameworks. The companion pieces cover Houston and Seattle–Tacoma.