S10 Group

Article - 11 MAY 2026

When Systems Stay Up but Care Starts to Slow

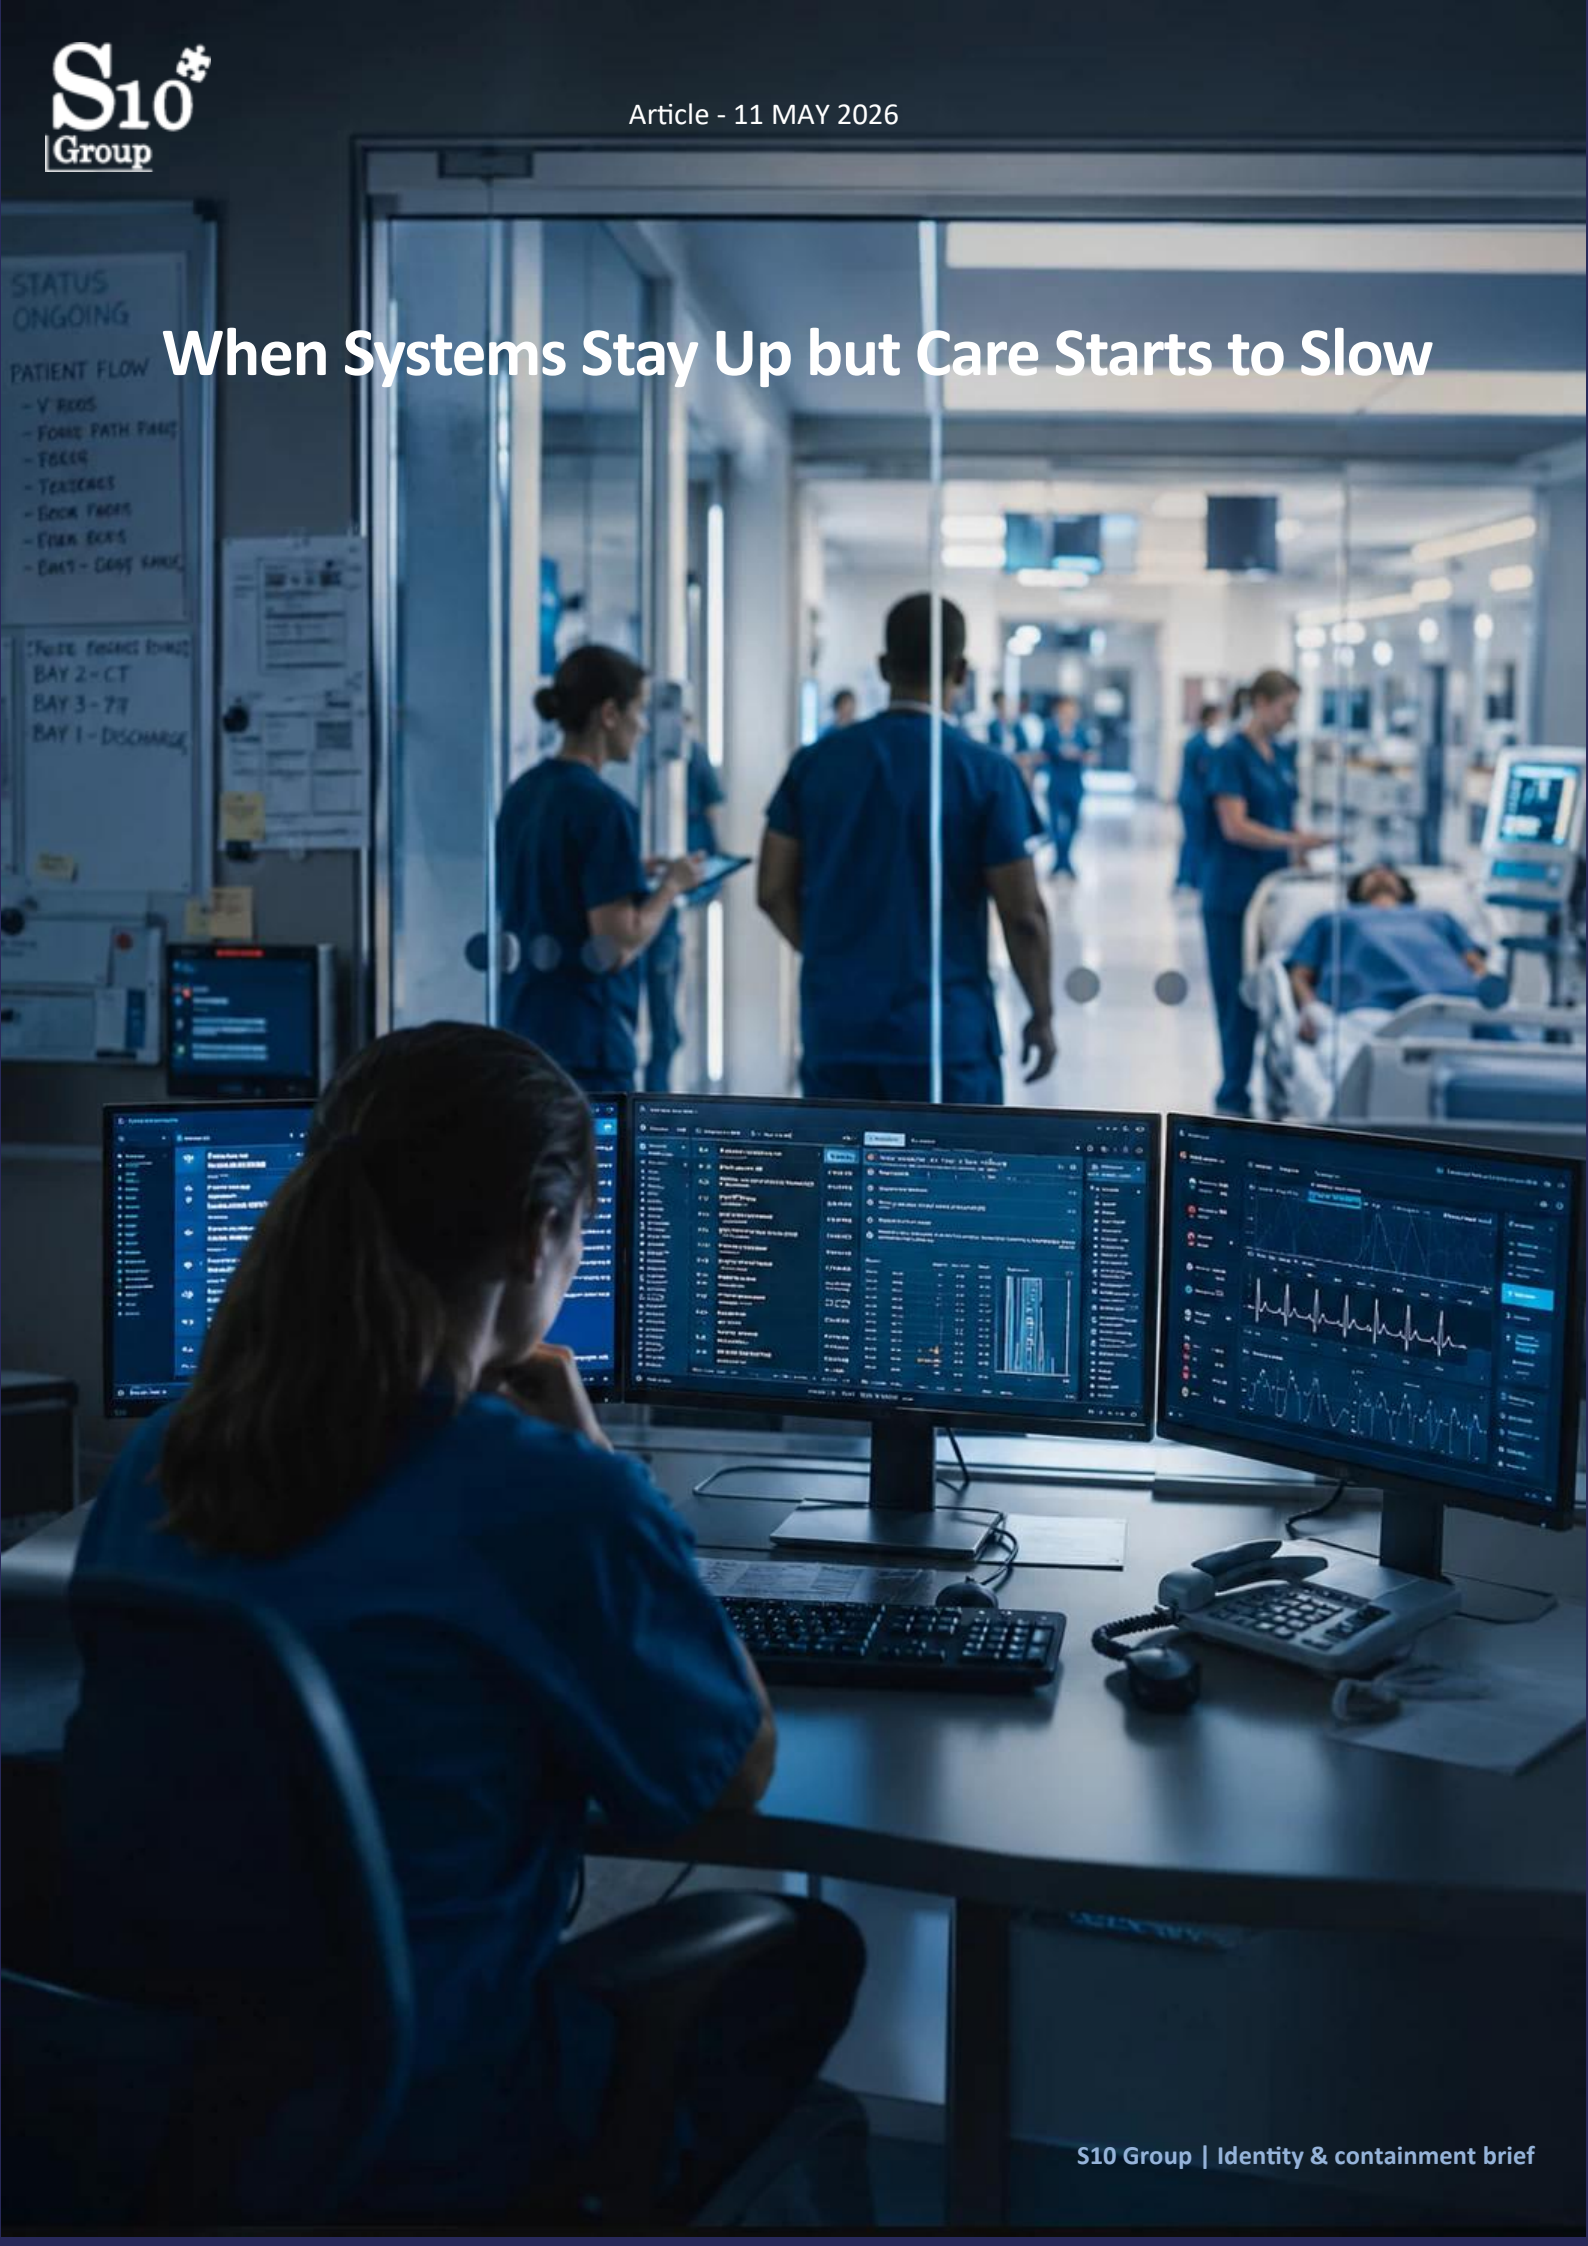

The first visible sign of cyber pressure in healthcare is not always collapse.

Sometimes the hospital is still open.

Staff are still logged in. Patient systems still respond. Wards continue to operate. Diagnostics are requested. Calls are answered. Patients are still being received.

And yet, something has already changed.

A result cannot be trusted quickly enough. A connection to a supplier becomes unsafe. A data exchange route is paused. A clinical team has to double-check what would normally be routine. A procedure is delayed because one dependency in the chain is no longer reliable.

That is often where the real incident begins.

Not when every system fails.

But when care starts to slow because digital trust has started to weaken.

The incident does not stay inside IT

Healthcare cybersecurity is often discussed through the language of systems, controls, compliance, and data protection.

Those things matter.

But they do not fully describe what happens when cyber pressure reaches care delivery.

Modern healthcare runs through digital trust. Patient records, pathology, imaging, medication workflows, referrals, monitoring, scheduling, billing, connected devices, external platforms, and communication between teams all depend on information being available and trustworthy at the moment of care.

When those systems are unavailable, the impact is obvious.

But when systems remain partly available while trust becomes uncertain, the situation is harder to govern.

Clinicians may hesitate. Administrative workarounds appear. Data has to be checked twice. External exchange may be paused. A routine workflow becomes slower because the organisation no longer knows which information, connection, or dependency can be relied on safely.

That is why a cyber incident does not stay inside IT.

It moves into operational continuity.

It moves into patient safety.

It moves into leadership decision-making.

- Pressure point

Still running is not the same as safe to trust.

Synnovis and the pathology dependency

The Synnovis incident in 2024 showed this clearly.

The attack was not simply a story about an isolated technology failure. It affected a central pathology dependency used by major London hospitals.

Blood tests, transfusion processes, appointments, and planned procedures were disrupted. Hospitals were still hospitals. Care did not stop everywhere. But a critical dependency became unreliable enough to change clinical decisions.

That distinction matters. The lights did not need to go out for care to slow.

If trusted pathology results, blood matching, or transfusion support cannot be relied on in time, the issue becomes clinical. Surgeons, clinicians, operations teams, and executives are forced into decisions that sit between cyber response and patient safety. The most important question is no longer only:

What system is affected?

It becomes:

What care decision now depends on a system or supplier we can no longer fully trust?

That is the healthcare version of control under pressure.

When one hospital is attacked, others feel it too

The wider healthcare system also shows another uncomfortable reality: the impact of a ransomware incident does not always stop at the organisation that was attacked.

Research on adjacent hospitals has shown that when one healthcare delivery organisation is disrupted by ransomware, neighbouring emergency departments can experience measurable strain. Patient volumes rise. Ambulance arrivals increase. Waiting-room times lengthen. Time-sensitive care becomes harder to deliver.

A separate study on cardiac arrest outcomes at untargeted nearby hospitals found that ransomware can create a spillover effect beyond the infected organisation itself.

This matters because healthcare resilience is not only an internal issue.

When one hospital loses capacity, patients move. When patients move, neighbouring hospitals absorb pressure. When neighbouring hospitals absorb pressure, regional care slows.

In other words: the cyber perimeter and the care perimeter are not the same thing.

A contained incident may remain local.

An uncontrolled incident can become regional.

- One lesson

In healthcare, resilience is not only about protecting one organisation. It is about preventing one organisation's loss of control from becoming pressure on the wider care system.

Patient safety is not an abstract consequence

This is why the human stakes cannot be treated as a side note.

The Alabama newborn case is often referenced because it shows the most difficult edge of the discussion. A lawsuit alleged that a ransomware attack contributed to a newborn's death after hospital systems were offline and staff could not see critical monitoring information in the normal way.

The details of individual legal cases are complex and should be treated carefully.

But the broader lesson is unavoidable.

Digital systems in healthcare are not administrative conveniences. They are part of how clinicians see, decide, prioritise, and intervene.

A monitor that does not show the right signal in time.

A lab result that cannot be trusted quickly enough.

A transfer that is delayed because a receiving hospital has absorbed diverted patients.

A planned procedure that is postponed because the supporting dependency is unsafe.

These are not only IT consequences. They are care consequences.

And they are exactly why cyber resilience in healthcare has to be measured by more than whether systems can eventually be restored.

Why compliance is not the same as control

Compliance creates a baseline.

It defines expectations, raises accountability, and improves discipline. It is necessary.

But compliance does not decide what happens when a live incident is already unfolding.

A policy does not isolate an unsafe dependency.

A framework does not decide which connection can be paused without unacceptable harm.

A completed assessment does not tell a hospital which workflow can continue in degraded mode when trust is unclear.

The gap is not between compliance and non-compliance.

The gap is between preparation and executable control.

That gap becomes visible when prevention has already been bypassed and the organisation must still decide what to do first.

- Board question

If a clinical dependency became unsafe tomorrow, who could decide what to restrict, what to keep running, and what level of degradation is acceptable to protect care?

Prevention can fail without the organisation being careless

This point matters because the wrong lesson from healthcare cyber incidents is often blame.

Healthcare organisations are not exposed because they do not care.

They are exposed because they operate complex, connected, time-critical environments where the margin for disruption is small.

Legacy systems remain in use because replacing them is difficult. Connected devices support care but expand the attack surface. External providers are necessary but create additional trust paths. Staff work under constant operational pressure. Clinical continuity often has to come first.

A compromised credential, an unsafe supplier connection, a missed signal, or a trusted pathway used in the wrong way can be enough for an attacker to move from the outside into the operational environment.

That does not make prevention irrelevant. It makes prevention incomplete.

Prevention reduces the chance of entry. It does not remove the need to control what happens if entry occurs.

What changes once something slips through

Once something slips through, the organisation enters a different phase.

The question is no longer only: How did this happen?

It becomes:

What can we still trust, and what must we restrict before the situation spreads?

In healthcare, that may mean temporarily limiting data exchange with an external platform, isolating a segment that shows abnormal behaviour, restricting privileged access, pausing a supplier route, or keeping a clinical workflow running in a narrower but safer mode.

In finance, the same logic may apply to transaction systems and privileged access.

In manufacturing, it may apply to production networks and remote maintenance paths.

In public services, it may apply to citizen-facing platforms and shared infrastructure.

The sectors differ.

The control question is the same.

Can the organisation still act while the facts are incomplete?

– Trust signal

Care can continue for a while on degraded systems. But it cannot continue safely for long on degraded trust.

Where impact is decided

A contained incident and a cascading disruption can begin in similar ways.

The difference often appears in the first decisions after trust becomes uncertain.

Can the organisation detect enough to know where pressure is forming?

Can it contain enough to limit spread?

Can it stabilise enough to keep essential operations moving?

Can leadership authorise action before the full report is available?

Those questions are not theoretical.

They decide whether an incident remains local, whether data exposure increases, whether care slows further, and whether recovery starts from a controlled position or from a wider breakdown.

This is why the first hour matters so much.

Not because everything can be known in the first hour.

But because the first hour often determines whether the organisation still has options.

A more realistic definition of resilience

Resilience is not the claim that every incident can be prevented.

It is the ability to keep operating when certainty is incomplete.

That means knowing which systems matter most, which connections can be narrowed, which identities can be restricted, which services can run in degraded mode, and who has authority to act immediately.

It also means recognising that care continuity depends on trust, not only availability.

A system that is online but unsafe to rely on may create a harder decision than a system that is clearly unavailable.

Because when something is clearly down, the organisation can switch to fallback.

When something is still running but uncertain, the organisation has to decide whether continuing to use it may create a larger risk.

That is the moment where leadership is tested.

– Operational reality

The first visible failure is not always the real beginning of the incident. In healthcare, the incident often begins when normal decisions start taking longer because certainty has already weakened.

Where S10 Group fits

S10 Group is positioned for the live phase after prevention has been bypassed and before the incident becomes much harder to govern.

The platform is designed to help detect malicious behaviour after entry, contain movement before it spreads further, reduce data-theft and ransomware leverage, and stabilise the environment while leadership still needs room to make decisions.

In healthcare, that room matters.

It can mean fewer systems becoming unsafe.

Fewer dependencies needing emergency restriction.

Fewer workflows slowing because trust has become unclear.

And a better chance that care can continue in a controlled, defensible way while the incident is still being understood.

The first question to pressure-test

If this happened tomorrow, what would you do first?
Would the organisation know what can still be trusted?
Would it know what can be safely restricted?
Would it know which clinical or operational functions must continue, even if parts of the environment are uncertain?
Would it know who has the authority to act before every fact is confirmed?
If those answers are unclear, the issue is not only technical.
It is a readiness gap.

And it will become visible precisely when the organisation has the least time to debate it.

The next question

This first article starts with the moment where care slows before systems fully fail.

The next article moves into the pressure behind that moment: ransomware in healthcare is no longer rare, isolated, or only about encryption. It is persistent, adaptive, and increasingly built around data, disruption, and leverage.

That is why the opening question cannot remain:

Are we compliant, or are we protected?

It has to become:

Can we stay in control when something slips through?

Control can still be regained — if a containment move exists.

- What containment changes

Containment does not remove the pressure of an incident. It changes how far that pressure can spread before leadership regains control.

SOURCES AND FURTHER READING

The article draws on public reporting and research into healthcare cyber incidents, regional care disruption, supplier dependency, and patient-safety consequences.

JAMA Network Open — Ransomware Attack Associated With Disruptions at Adjacent Emergency Departments in the US

Research on how a month-long ransomware attack affected adjacent emergency departments, including patient volume, ambulance arrivals, waiting times, patients leaving without being seen, and acute stroke-care pressure.

NHS England — Synnovis ransomware cyber-attack

Operational updates on the Synnovis incident and its effect on pathology services, appointments, blood tests, and care disruption across affected London hospitals.

Healthcare IT News — Hospital ransomware attack led to infant's death, lawsuit alleges
Reporting on the Alabama newborn case, framed carefully as a lawsuit allegation, illustrating why patient-safety consequences must be discussed with care