CSL

Critical IoT Connectivity **Experts**

Healthcare Connectivity Architecture

Clinical Alarm and Telemetry Pathways and the NHS Connectivity Standards



Secondary Care | Briefing

Meeting NHS England's networks and connectivity standards does not by itself separate a hospital's remote cardiac alarms from its patient record. If their shared path fails, staff usually notice they cannot open the patient record; a lost alarm may go unnoticed. Five questions test each alarm pathway as well as the site.

For CIOs, CCIOs, network and connectivity leads, clinical safety officers, clinical engineering, estates and procurement leads

Format
Briefing

Date
October 2026

21 standards | Five questions | A record for each pathway

Contents

In brief 3

1. One hospital, several networks, one upstream path 4

2. NHS England's 21 standards are written for sites, not clinical alarm and telemetry pathways 6

3. Five questions for each clinical alarm or telemetry pathway 8

4. From self-assessment to plan 11

5. CSL Group products by site 14

References 15

Figures

Figure 1: One site edge for every pathway, beside the same site with its clinical alarm and telemetry pathways separated 8

Figure 2: Several radio networks reaching one operator core, and two independent operator cores 9

Tables

Table 1: What happens in each domain when a shared upstream path fails 5

Table 2: What each group of standards addresses at a site 7

Table 3: What a record kept with the self-assessment can hold 13

Table 4: Upstream paths by device or site 14

In brief

On 13 August 2026 NHS England published 21 networks and connectivity standards for every NHS organisation in England[1]. Each organisation is asked to assess itself against them every year. Where a standard cannot yet be met, the organisation is expected to keep a plan setting out the risks, the interim mitigations and a timeline[1]. The first assessment will therefore record each site's position against the standards, and what remains to be done.

The standards set specific, measurable expectations for each site's network. Each site should have a backup connection designed for supplier, vendor, technology and, where possible, route diversity[2]. Failover to it should take no more than 10 seconds and should be tested every quarter[3]. Core network equipment should be powered through an uninterruptible power supply (UPS)[4], and traffic should be segmented by user group and device type[5]. Where care depends on Wi-Fi, the network "must be designed, assured and monitored as a clinical service"[6]. In standards 1.2, 3.1, 5.2 and 8.6, the level of resilience follows clinical risk[2,5,6,7].

The standards apply that principle to each site and name its critical areas and high-risk devices. Because they are written site by site and layer by layer, medical devices appear in them as systems that depend on the network, not as pathways whose alarms need a route of their own. None of the 21 standards names a lift alarm, nurse call or the Health and Social Care Network (HSCN). Building management systems appear only once, in standard 2.2's list of what a loss of power to core equipment "can immediately disrupt access to"[4]. Nor do the standards name alarm pathways: the routes that carry a medical device alarm, a telemetry feed, a nurse call or a lift alarm call to the people who respond.

Remote cardiac alarms, building management traffic and electronic patient record (EPR) traffic can still leave a hospital through the same site edge and perimeter. Duplicated routers or firewalls make that edge more resilient, but they do not give any of this traffic a route of its own. A record kept with the annual self-assessment can show which pathways share that edge, provided it counts pathways as well as sites. Fire and security alarm signalling is outside this briefing's scope, although CSL's alarm signalling products and services carry those signals to alarm receiving centres (ARCs) and CSL also provides connectivity services directly to ARCs (section 5). Specialist installers, including CSL Group's accredited installer partners and large system integrators, fit that signalling equipment.

This briefing is organised around five questions that a Trust can ask of each alarm or telemetry pathway it identifies. Section 3 takes them in turn, with the standards bearing on each:

1. what the pathway shares with clinical IT on the site;
2. what it shares in the comms room;
3. whether its second path is independent;
4. whether a lost alarm would be noticed;
5. who owns it.

Four rules follow from the five questions:

1. Count the dependencies a pathway shares with clinical IT, from the core switches to the clinical wide area network (WAN), not only the number of suppliers it uses. A pathway bought from a different supplier can still depend on the site's router and firewall, or on the core switches and on the comms room's power and cooling. Each shared dependency is a common failure domain: a fault there can affect the alarm and the clinical systems together.
2. Match the independence of each pathway's upstream path to the consequence of losing it, which applies the standards' own principle to pathways as well as sites.
3. Test a claim of independence at the operator core as well as the access network, because a backup over one mobile operator diversifies the technology and leaves one core in the path.
4. Ask the receiving platform for an acknowledgement or delivery record for each alarm, and name who is alerted when an alarm is not acknowledged or recorded.

Where a ward local area network (LAN) and its Wi-Fi are otherwise suitable, they can stay as they are. Under standard 5.1, Wi-Fi at NHS sites should be delivered "over a shared physical infrastructure wherever practicable, using logical separation and policy-based controls"[6]. Where a pathway leaves the site and its loss would justify it, as with remote cardiac alarms, infusion alerting and lift alarm calls, it can be given a path of its own. CSL provides that path in the device, at a gateway or at the site, without replacing the ward network (section 5). Managed print, staff Wi-Fi and similar services can keep sharing the upstream path once the risk is recorded and accepted.

1. One hospital, several networks, one upstream path

A hospital runs several operational domains rather than one network. Clinical IT, medical technology, life-safety signalling and estates systems each have their own owner, assurance route and consequence of failure. Yet they can leave the site through one upstream path, and each Trust needs to establish which of its alarm and telemetry pathways depend on it. In England that path often reaches HSCN, to which many NHS organisations have connected[8].

Each domain has its own owner, assurance route, consequences and recovery route

- **Different owners.** Trust IT, clinical engineering and estates are separate teams.
- **Different assurance routes.** Clinical safety, cyber security and life-safety signalling are each assured against their own standards, by people with different operating responsibilities.
- **Different consequences.** A failed EPR, or a lost connection to the imaging archive (PACS), is generally noticed by the staff using it. A cardiac alarm that does not reach the central station, by contrast, may raise no error there unless the connection is supervised. Even then, the station shows that the link has failed, not which alarms were lost (section 3).
- **Different recovery routes.** Trust IT restores clinical IT systems, such as the EPR, through its incident management process. Clinical engineering restores medical technology, including cardiac monitoring, while estates restores the lift alarms and the building management and water safety monitoring systems. Each team follows its own procedures and sets its own priorities.

A shared outage becomes several incidents at once

Where several domains share an upstream path, one failure can start concurrent incidents in services that are otherwise separate. Telemetry visibility, clinical alarm transmission, lift alarm calls, remote plant reporting and access to clinical platforms can then be lost together. Where those services depend on the external path, a cardiac alarm does not reach a central station hosted off site and a lift alarm call cannot connect. Alarm pathways leave the site where the central station or platform is hosted elsewhere, where one hospital's telemetry is watched from another, or where a lift alarm call goes to an ARC. Where the central station runs on site, an alarm depends first on equipment inside the hospital: the ward network or the telemetry receivers, the monitoring servers and, where alarms go on to staff handsets, the alarm server that relays them. The questions in section 3 on the comms room, on delivery and on ownership apply to those pathways as well.

The shared path is also a cyber dependency. If ransomware compromises the Trust network, an alarm or telemetry pathway with an upstream path of its own may remain available, but only where its routing, authentication, monitoring and downstream dependencies are maintained independently of the Trust WAN.

A site can fail over in 10 seconds with both connections behind one firewall

Standard 1.2 asks each site for a diverse backup connection, supported by multiple routers or router modules, appropriate routing configuration and UPS protection[2]. Its 10-second measure applies to the switch to the backup connection[3]. Even where those measures are met, an alarm pathway can still depend on equipment both connections pass through, such as the firewall and the core switches. It can also depend on the comms room's power and cooling, and on a network that may be isolated during a cyber incident. In those failures the EPR and the remote cardiac alarms are lost together. Each alarm pathway therefore needs its shared dependencies identified and tested, which section 3 sets out.

Who restores a shared path depends on what failed

If a circuit or router fails, Trust IT and its network supplier restore it. If the comms room loses power or cooling, estates or IT restores them first, and the network comes back afterwards. Until then, the teams whose services depend on the path, such as clinical engineering for remote cardiac alarms, cannot restore it themselves, because the fault is in equipment they neither own nor repair. Table 1 sets out the incident each domain sees and the team that owns it.

1. One hospital, several networks, one upstream path (continued)

Domain	What the shared path carries	What happens when it fails	Who owns the consequence
Clinical IT	EPR, the patient administration system (PAS) and PACS hosted off site, and remote radiology reporting	Staff generally notice the loss of access and follow the Trust's downtime procedures	Trust IT, through its incident management process
Medical technology	Remote cardiac alarms to a hosted central station, and infusion alerts to a platform off site	Where the device has bedside alarms, they continue, but alarms no longer reach the central station or remote platform. A supervised connection reports the failure; an acknowledgement or delivery record shows which alarms arrived	Clinical engineering
Life-safety signalling	Lift alarm calls carried over the Trust network	A lift alarm call cannot connect	Estates
Estates systems	Building management and water safety monitoring hosted off site	Remote plant reporting stops; whether estates receives a loss-of-communications warning depends on the monitoring arrangements	Estates

Table 1: What happens in each domain when a shared upstream path fails, and which team owns the consequence.

In recent NHS incidents, one failure reached many services at once

In the first of two incidents, many services depended on one pathology supplier, which served several NHS organisations. In the second, failed cooling equipment affected both digital systems and critical clinical services at one hospital. In June 2024 a ransomware attack on the pathology supplier "significantly reduced its capacity to process tests", with the greatest impact in south-east London, according to NHS England[9]. By late September NHS England's London region had reported provisional figures for the two most affected Trusts of 10,152 postponed acute outpatient appointments and 1,710 postponed elective procedures[9].

During the heatwave of June 2026, Portsmouth Hospitals University NHS Trust declared a critical incident at Queen Alexandra Hospital after "the failure of several chiller units supporting critical infrastructure"[10]. The Trust reported that the loss of cooling had raised temperatures in a number of areas, "affecting our digital systems and critical clinical services". It stood the incident down three days later[10].

Neither report describes an alarm or telemetry pathway, so they do not show that one failed. An alarm or telemetry pathway can be exposed to both kinds of failure: to the first where one supplier hosts its central station or alarm platform for several Trusts, and to the second where its equipment sits in a room whose chillers can fail. Its loss might also go unnoticed (section 3).

Beyond these two incidents, the same pattern of one failure reaching many services appears in a larger set of documented failures. CSL's Critical Connectivity® Failure Atlas analyses 45 publicly documented infrastructure failures from 2018 to 2026, in areas that include telecommunications networks, emergency calling, cloud and identity services, and data centres[11]. The Atlas selects consequential, well-documented failures whose spread can be traced, so every one involved a dependency that several services shared. Its counts therefore describe those incidents, not how often failures occur. No alternative was recorded as preserving the complete intended service, and an alternative path carried partial traffic in six. In 27 of the 45, the monitoring, management access, credentials, configuration data or communications needed to restore the service were themselves affected, so recovery was impaired along with the service. Healthcare services were affected in seven.

The Atlas also cites a study of a 2024 endpoint-security software update. After it, internet-facing services stopped responding at 759 of 2,232 US hospitals. That study covered hospitals that used the same electronic health record product and had internet services that could be observed from outside[11].

2. NHS England's 21 standards are written for sites, not clinical alarm and telemetry pathways

Standards 1.2, 3.1, 5.2 and 8.6 scale resilience to clinical risk, and they apply that scaling to each site and its network as a whole rather than to each alarm or telemetry pathway.

Twenty-one standards in eight groups, assessed every year

The 21 standards set out how NHS organisations should build, manage and secure their networks[1]. They were produced by NHS England's Networks and Connectivity programme, which is part of its Frontline Productivity portfolio. They are published by the programme's Networks and Connectivity Centre of Excellence[12]. The Centre states that it exists "to define, promote, and uphold national standards" for secure, resilient and high-quality network and connectivity infrastructure[12].

The standards apply to every NHS organisation in England, and "should be considered applicable to all NHS organisation, site and network types"[1]. Their scope is therefore every site and every kind of network a Trust runs, not only the clinical network. The 21 standards are arranged in eight groups: fixed line connectivity, network switching, network cabling, network management, wireless at NHS sites, wireless on the move, leadership and governance, and security and compliance. Each standard sets out the importance of meeting it, when and how to meet it, and, in all but one, its technical requirements[1,13].

An organisation is expected to assess itself every year and whenever it makes major changes to its network and connectivity estate[1]. A further assessment follows organisational changes or mergers that affect ownership of, or responsibility for, sites or infrastructure. Where a standard cannot yet be met, the organisation should keep a plan naming the risks and impacts, the interim mitigations and an anticipated timeline[1]; section 4 shows how the same plan can record a shared upstream path.

Every standard is headed "should", although "must" appears in parts of the supporting text and the standards page calls them "ongoing requirements"[1]. This briefing therefore describes them as NHS England's expectations, and quotes "must" where the text uses it.

The four standards set resilience by clinical risk, site type and dependence on the network

The backup connection standard states that "the specific level of redundancy required will depend on" clinical risk, operational complexity, site type and size, and business continuity needs[2]. The cabling standard provides for additional resilience "in proportion to clinical risk" where clinical services depend on the network[7]. Where care depends on Wi-Fi, enhanced capability "must be provided proportionate to clinical risk"[6]. The perimeter standard includes failover paths "where clinically required"[5], which keep the perimeter itself available rather than give a pathway a route around it.

CSL's white paper Resilient Connectivity for Secondary Care (April 2026) applies the principle to each connection: a connection should be as resilient as the clinical consequence of losing it requires. These four standards apply it at the level of the site, and standards 5.3 and 8.4 single out critical areas and high-risk devices within it[5,6]. None of the 21 standards, however, names an individual alarm or telemetry pathway, such as a cardiac monitor's alarm to the central station or a lift alarm call. A site that meets the standards can therefore still send its remote cardiac alarms, where the central station is hosted off site, out through the same site edge and WAN as its EPR. Section 3 applies NHS England's principle one level further down, to each pathway.



2. NHS England's 21 standards are written for sites, not clinical alarm and telemetry pathways (continued)

Meeting a group's standards still leaves shared equipment, delivery and ownership to check

Each group of standards addresses particular failures at a site, such as the loss of its primary connection or of power to its core switches. Once a group's standards are met, an alarm or telemetry pathway can still be exposed. It may depend on equipment it shares with clinical IT, such as the site's firewall or the core switches in the comms room. A failure of that equipment can then stop the alarm as well as the clinical systems. An alarm may also be lost with nothing at its destination to show it, or its pathway may be owned by estates or clinical engineering and fall outside the network owner's plans. For each group, Table 2 gives the failures its standards address and, in its right-hand column, what can still leave a pathway exposed once they are met.

Group of standards	What the standards address	What still leaves an alarm or telemetry pathway exposed
Fixed line connectivity (1.1, 1.2)	Loss of the site's primary connection: a diverse backup with failover within 10 seconds, tested quarterly, multiple routers or router modules, and UPS protection	The firewall and other equipment at the site edge that both connections pass through, and the network beyond them. A mobile backup through one operator also depends on its core
Switching and cabling (2.1, 2.2, 3.1)	Supported switching equipment, resilient uplink connections and UPS protection for core equipment	The core switches, and the comms room's power and cooling. Where clinical systems, Wi-Fi, telephony and building management share those switches, one power or cooling failure in that room can stop the alarms along with those systems
Network management (4.1 to 4.5)	Undetected degradation and unsupported equipment: monitoring, central management and software-defined networking (SD-Networking)	Monitoring that watches the network's equipment but cannot show whether each alarm reached its destination, so an alarm that is lost may not be noticed
Wireless at NHS sites (5.1 to 5.3)	Poor coverage and reliance on Wi-Fi alone: surveyed Wi-Fi and indoor cellular in critical areas	The site's upstream path, which a device on Wi-Fi still needs to reach a platform off site. Good Wi-Fi coverage on the ward does not make that path independent
Wireless on the move (6.1)	Staff working off site without connectivity[1]	Public mobile networks. Devices that travel with patients or staff depend on them. With one operator, they also depend on its core, however many radio networks they can use
Leadership and governance (7.1)	Networks with no owner, and continuity plans that leave connectivity out	Pathways that estates or clinical engineering own, which can fall outside the network owner's plans and failover tests
Security and compliance (8.1 to 8.6)	Unauthorised access, flat networks and attack at the perimeter: access control, segmentation and perimeter security	Under standard 8.4, traffic is separated into segments by user group and device type. Segmentation keeps one segment's traffic apart from another's, which limits how far a compromise or a fault in one segment can spread. Where the segments share switches, cabling, links and a perimeter firewall, segmentation does not protect against a failure of that equipment, which can stop every segment at once

Table 2: What each group of standards addresses at a site, and what can still leave an alarm or telemetry pathway exposed once the group's standards are met. The switching and cabling groups share a row, because both concern the core switches and the cabling to them. The right-hand column describes a site where the pathways share one edge and perimeter, as section 3 sets out, and is CSL's description of what each group leaves shared, not NHS England's.



3. Five questions for each clinical alarm or telemetry pathway

Each question below covers one stage of a pathway's route, from the site edge to the team that owns the alarm, and identifies the NHS England standards relevant to assessing it.

What does the pathway share with clinical IT on the site?

The pathways that use the site network share its edge. Remote cardiac alarms bound for a central station hosted off site, the building management system where it shares the clinical WAN, and EPR traffic all leave through the same router and firewall, whichever circuit is live. Where that edge or the network beyond it fails, or is isolated during a cyber incident, those pathways lose both circuits at once.

The Failure Atlas described in section 1 draws the principle behind this exposure from its 45 incidents: "Redundancy protects only against the failure domains from which it is independent"[11]. A backup connection is independent of the primary circuit, but not of the router and firewall that both connections pass through. Figure 1 shows such a site beside the same site with its alarm pathways given their own upstream paths.

Standard 8.4 states that clear separation of traffic "helps protect sensitive systems while allowing shared use of network infrastructure"[5]. Where segmented services share the site edge, the perimeter and the upstream provider, however, logical segmentation does not remove those dependencies.

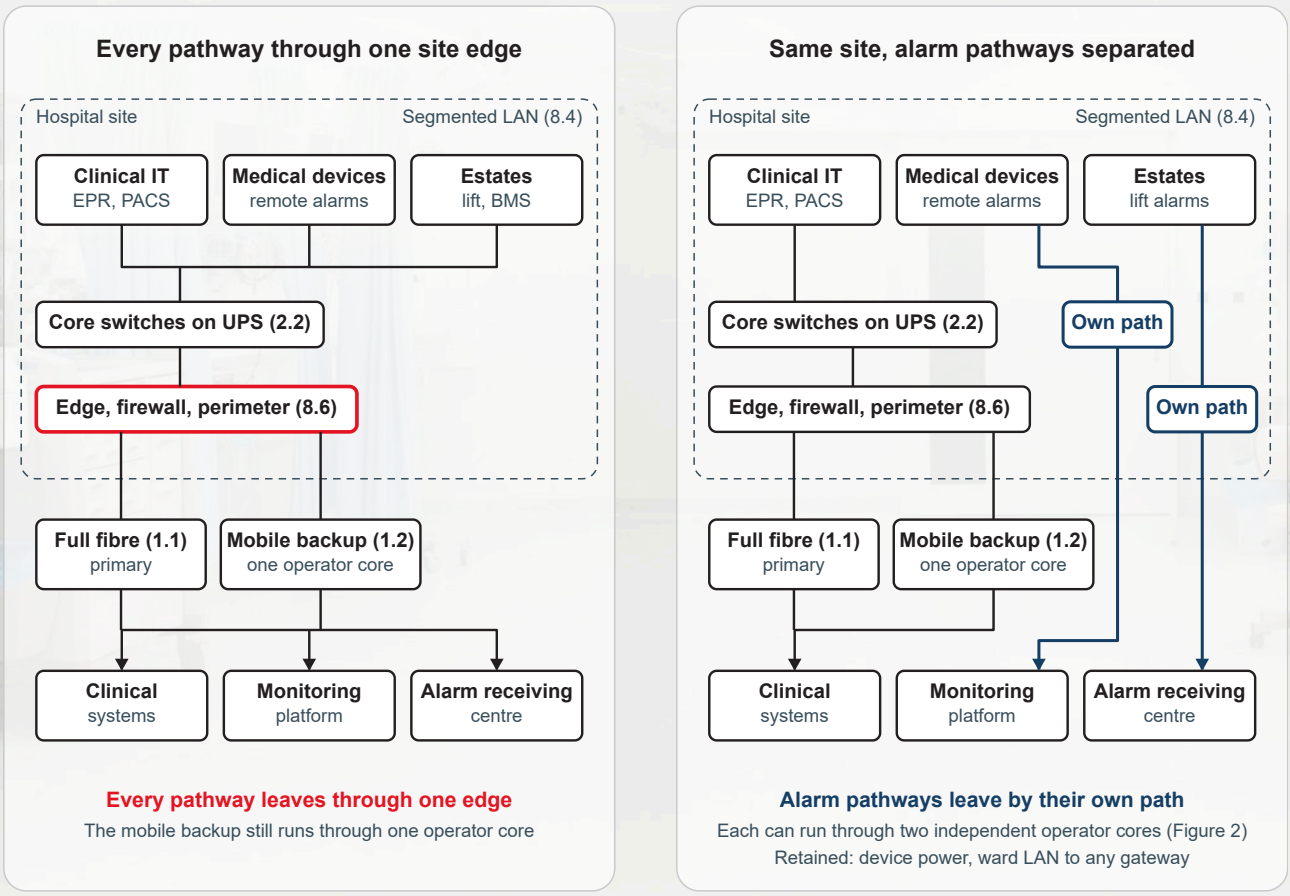


Figure 1: One site edge for every pathway, beside the same site with its clinical alarm and telemetry pathways separated. Both views show a diverse backup, UPS, segmentation and perimeter security. In the separated view, remote medical alarms and lift alarm calls leave by upstream paths of their own, each of which can run through two independent operator cores (Figure 2). Building management systems, not drawn there, stay on the core switches. Where present, shared gateway power, backhaul and receiving-platform dependencies need assessing separately.

3. Five questions for each clinical alarm or telemetry pathway (continued)

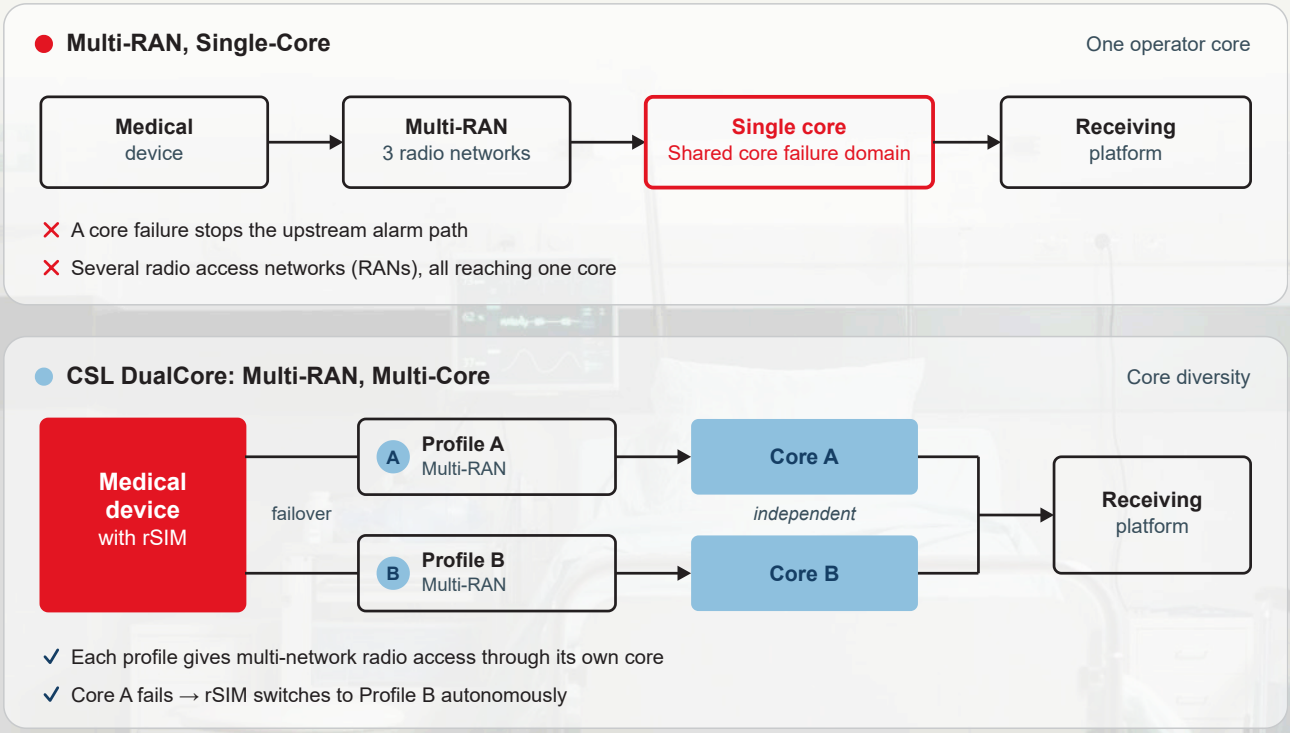


Figure 2: Several radio networks reaching one operator core, and two independent operator cores. Above, a SIM with one operator profile can use several radio networks and still depend on that operator's core; below, rSIM holds two independent multi-network profiles on one SIM and switches to the second, on an independent core, if the first core fails.

What does the pathway share in the comms room?

An alarm pathway that leaves the site through the comms room depends on the room's power and cooling, and on any core switches it shares with other systems. Standard 2.2 lists building management systems beside clinical systems, wireless networks and telephony among the systems that a loss of power to core equipment "can immediately disrupt access to"[4].

A UPS provides "short-term battery backup"[4], so every pathway through the comms room depends on its runtime until standby generation takes over. If generation fails, the same runtime sets how long those pathways survive in total. Cooling is a dependency of the same kind: at Portsmouth in June 2026, the failure of chillers supporting critical infrastructure affected digital systems and critical clinical services together (section 1)[10].

Is the pathway's second path independent?

Standard 1.2 gives examples of diverse primary and backup combinations, such as full fibre (FTTP) with low Earth orbit satellite, or FTTP with 4G and 5G mobile[2]. Where the backup is mobile and uses one operator profile, it still reaches the wider network through that operator's core, however many radio networks it can use. A fault in that core removes the backup whatever the strength of the radio signal on site, as Figure 2 shows. NHS England's own guidance on mobile backup, written mainly for small sites such as GP surgeries and care homes, advises choosing "different network providers for different SIMs" where resilience is the aim[14].

The same guidance asks that a virtual operator's SIM be checked for a different underlying network, and notes that "the trigger for switching is usually a total loss of connection"[14], so a connection that is degraded but not lost may not trigger a switch. For a second path that also survives a fault in an operator's core, CSL builds both paths to three design principles: independent operator cores, switching made by the SIM, the device or the router itself, and private routing on both paths. It also supplies both paths: rSIM® carries them on one SIM, and a CSL DualSIM® Router configured for core independence uses two SIMs, both supplied by CSL. In the configuration CSL specifies, each path runs through its own operator's core and then over CSL's private access point name (APN) routing, with an encrypted connection to the receiving platform.

3. Five questions for each clinical alarm or telemetry pathway (continued)

For a site's EPR and administrative traffic, a single-operator mobile backup may be proportionate. A pathway whose loss could delay the response to a cardiac arrest can justify a backup that survives a fault in one operator's core. For such a pathway, CSL recommends testing the supplier diversity set out in standard 1.2 at the operator core as well as at the access network.

Standard 1.2's 10-second measure (section 1) applies to the site's switch to its backup connection[3]; for an alarm pathway, the device or router must first detect the failure, and what it can detect depends on what has failed. rSIM holds two independent multi-network roaming profiles on one SIM, so it can respond whether a radio network or an operator core has failed. Within a profile, a device that loses one radio network can move to another as soon as it detects the loss. A core outage can leave the radio signal intact, and a SIM with one operator profile has no second core to move to. Logic on the rSIM monitors the active connectivity path, and once it detects a sustained loss of usable connectivity, it switches to the second profile, on an independent core, without waiting for an instruction from a platform. A specification can then state the time each alarm pathway allows, and ask for that time to be shown in a core failover test.

In the incidents the Failure Atlas analyses, a second path fell short in two ways. Either it shared the failure domain of the first, or it lacked the capacity or functionality the event demanded[11]. In the 2023 999 emergency call incident, Ofcom found that the disaster-recovery platform lacked the capacity and functionality needed for reasonably expected demand[11]. A recovery arrangement existed, yet the service was still not preserved. For an alarm pathway, a failover test is therefore most informative when it runs at realistic load, scope and duration, and when it records the outcome observed rather than the intended outcome[11].

Would a lost clinical alarm be noticed?

Standard 4.3 asks for monitoring of network performance, capacity and health, including latency, jitter and utilisation, and for end-user experience monitoring of clinical applications (Table 2)[1]. An alarm pathway has to be judged where the alarm arrives: the test is whether the alarm or call reached the people who respond within the time the clinical or life-safety process allows. Neither kind of monitoring therefore shows whether an individual cardiac alarm reached the central station within the required window.

A WAN that passes a reachability test may still lose enough packets, or add enough delay, to make it unusable for time-sensitive alarm and telemetry traffic. Packet loss and delay are not the only faults that such a test can miss. In 25 of the 45 failures in the Failure Atlas, the initiating fault lay above the bearer, in a platform or service that traffic reaches rather than in a carrier's network, and a device can keep its radio signal while such a service is unavailable[11]. Monitoring limited to signal and reachability can therefore report a healthy connection while another part of the service has failed.

Where a pathway is supervised end to end, a failure to deliver raises its own alert. Supervision of the connection alone, such as a periodic heartbeat, shows that the path is up; showing that an individual alarm arrived needs an acknowledgement or delivery record for that alarm from the receiving platform. Where there is no such record, the first sign of a lost alarm may be the absence of a response. What a device shows locally when its connection is lost, and what the central system records or recovers afterwards, depend on the device's intended functions and on how the system is configured. They are therefore established for each installation rather than assumed. A record kept with the self-assessment can note them for each pathway, alongside whether delivery is supervised and who receives the alert.

Who owns the pathway?

Under standard 7.1, each organisation should assign a member of its senior leadership team to networks and connectivity[13]. That lead should have strategic oversight and should ensure that the organisation's position against the standards "is reviewed annually"[13]. The lead should also ensure that "networks and connectivity requirements are included in disaster recovery and business continuity plans"[13]. Under the standard, the lead is assigned at the first assessment, and again when the lead leaves or changes role or when responsibilities in the leadership team change[13].

The standard gives the network an owner, whereas an alarm pathway usually has several. Clinical engineering owns the device, and estates the lift alarm. The contract for the receiving platform or the ARC sits with the team that buys it, such as clinical engineering or estates. The path through the site network belongs to Trust IT. The governance group for a pathway can therefore include the chief information officer (CIO), the chief clinical information officer (CCIO), the clinical safety officer, clinical engineering, estates and cyber security. The record kept with a Trust's self-assessment can name the owner of each pathway alongside the lead for the network.

4. From self-assessment to plan

Counting pathways as well as sites shows which alarm pathways still share the upstream path. The plan kept under the standards can then record the decision for each, and its date.

A shared pathway can be recorded in the plan

The plan described in section 2 records the risks, the interim mitigations and a timeline for each standard not yet met[1]. A pathway whose upstream path is shared with clinical IT can be recorded in the same form, even where every standard is met:

- the risk is the loss of the pathway when the shared path fails;
- the interim mitigation is the procedure staff follow while the path remains shared;
- the timeline is the date by which an independent path, or an explicit acceptance, is in place.

Six steps from one pathway to the estate

1. Identify one high-risk pathway or asset class, such as remote cardiac alarms, infusion alerting or lift alarm calls.

2. Map its path from the device to its destination, and record each switch, circuit, platform and service it shares with clinical IT.

3. Confirm the governance owners for the pathway across the CIO, CCIO, clinical safety officer, clinical engineering, estates and cyber security.

4. Decide whether separation is required, desirable or unnecessary, based on the consequence of loss and the time to recover.

5. Pilot one pathway, and test failover in a controlled window, with downtime procedures in place, before relying on it.

6. Scale only after clinical safety and operational sign-off.

Three classes of pathway to assess first

Each of these classes combines a high consequence of loss with a remote alarm that can depend on the shared path:

- high-acuity monitoring and alarm workflows, such as cardiac telemetry and vital signs;
- infusion and drug delivery, including patient-controlled analgesia and syringe drivers that send alerts off site;
- life-safety signalling and critical estates monitoring, including lift alarm calls.

Services such as managed print may share the path once the risk is accepted

Outside those classes, a shared dependency may be justified for bedside entertainment, managed print, staff Wi-Fi, non-critical building sensors and supplier-managed catering systems. The condition is that the clinical and operational risk is understood, documented and accepted.

Each team buys connectivity under contracts of its own

Three of the owners named in section 3 hold the contracts under which connectivity is bought:

- Trust IT buys the WAN and network services, and may manage the mobile contract as well;
- estates may buy the lines for lift alarms inside a maintenance contract, or as a small purchase of its own;
- a monitoring platform and its connectivity may come under the device manufacturer's service contract, which clinical engineering manages.

These contracts may not record dependencies pathway by pathway. Whether a pathway shares the upstream path is decided by the architecture as deployed, and a Trust can state the architecture it needs in the specification for the contract that covers the pathway.

Four of the national agreements under which Trusts buy network and mobile services change between September 2026 and July 2027. For mobile services, Mobile Voice and Data Services (RM6261) ends on 12 December 2026, and Mobile Voice and Data Services 2 (RM6391) launched on 17 September 2026. Network Services 3 (RM6116) runs to 17 July 2027, with its lots 1c and 1d ending on 22 May 2027. The Government Commercial Agency, formerly Crown Commercial Service, expects to award its replacement, Network Services 4 (RM6377), in March or April 2027. The product types listed for Network Services 4 include IoT and connected places, paging and alerting services, and critical communications, and its lot structure is expected to differ from Network Services 3's[15]. HSCN is bought under an agreement of its own (RM3825), which runs to 23 May 2028. A Trust that buys under one of the new agreements can write into its specification the pathways that need an upstream path of their own.

4. From self-assessment to plan (continued)

Each Openreach analogue line migrated by 31 January 2027 is a choice of path

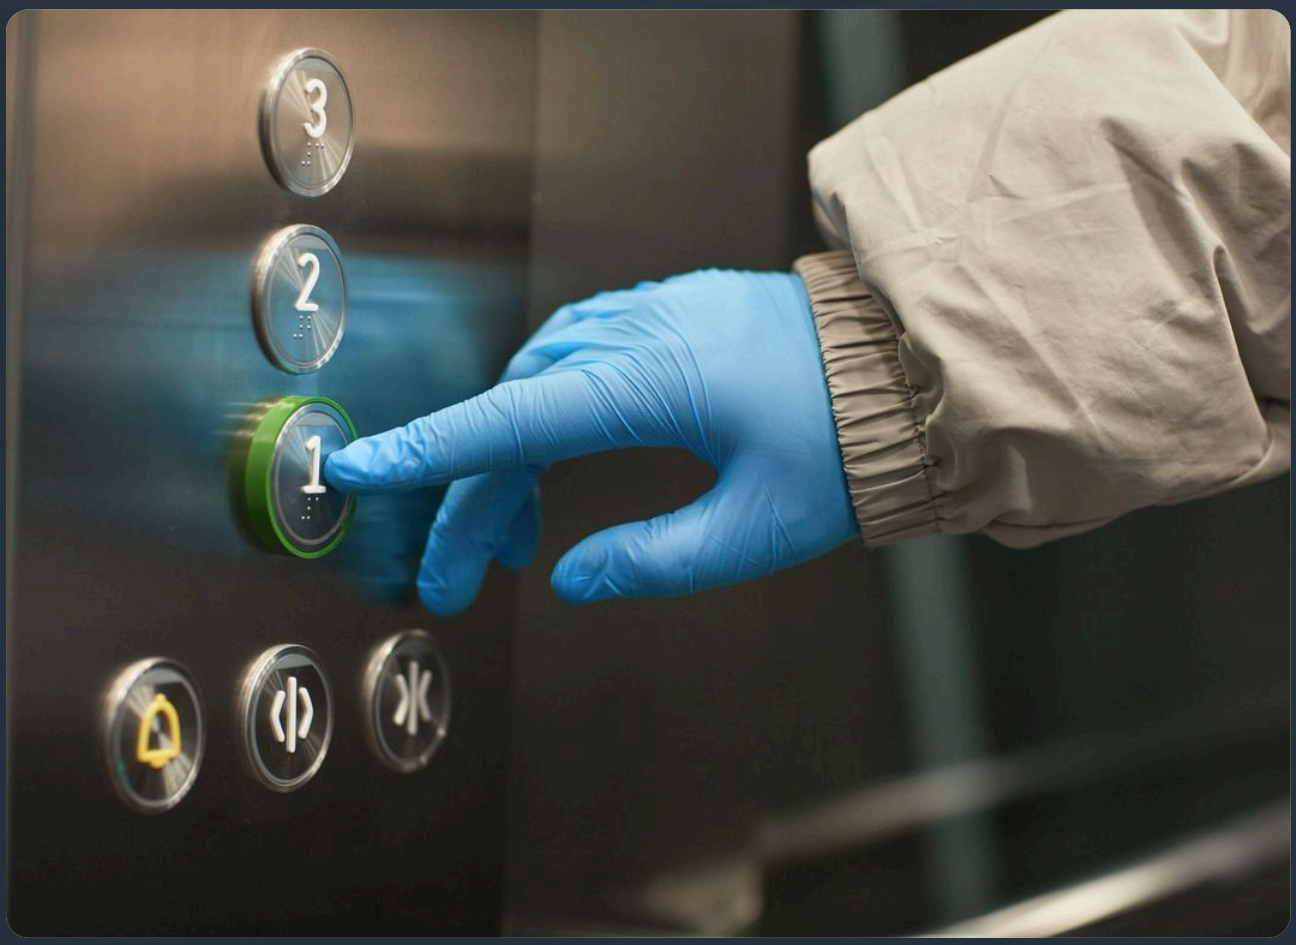
For lift alarms that still signal over an Openreach analogue telephone line, the date for moving them is already fixed. The Department for Science, Innovation and Technology (DSIT) publishes the government's timeline, which states that all users of the Openreach analogue network "will need to be migrated to new services by the 31 January 2027"[16]. DSIT's landline guidance, in its section for businesses, names "lifts and elevators" among the systems affected[17]. The department's charter for critical national infrastructure names health among the sectors where loss of PSTN or ISDN connectivity "could result in loss of life or casualties"[18].

UK mobile network operators plan to switch off 2G between 2029 and 2033, and DSIT's 2G guidance names lift alarms among the equipment that can rely on it[19].

A lift alarm moved onto the Trust network depends on that network's shared edge and comms room. One moved onto a cellular unit that relies on 2G alone would need replacing again when the 2G networks close. One moved onto one operator's service depends on that operator's core, however many radio networks its SIM can use. Recording the choice with the self-assessment keeps it visible after the analogue line has gone.

One record for each pathway, kept with the self-assessment

The six steps build one record for each pathway, and Table 3 sets out what it can hold. Its entries follow the five questions in section 3, with the plan's decision and the contract added from this section. The record then shows in one place what a pathway shares, how its loss would be noticed, who owns it and how it is bought. A Trust can keep it in the form its self-assessment and plan already use. NHS England's guidance on connected medical devices already asks for a drawing of how each device communicates with other devices and services[20], and the map in step 2 can start from that work. The manufacturer may be able to supply documentation for it, such as a Manufacturer Disclosure Statement for Medical Device Security (MDS2) form[20].



4. From self-assessment to plan (continued)

Entry	What the record holds for the pathway	Where this briefing covers it
The pathway	The device or system that raises the alarm, where the alarm goes, on site or off site, and the people who respond	Sections 1 and 3
Shared with clinical IT on the site	The site router, firewall and perimeter the pathway leaves through, and whether the clinical WAN carries it	Section 3, first question; standards 8.4 and 8.6
Shared in the comms room	The core switches it passes through, the UPS and its runtime, and the comms room's cooling	Section 3, second question; standard 2.2
The second path	The access network and operator core of each path, the time the pathway allows for a switch, and the date and result of the last core failover test; for a pathway that stays on site, the resilience of its local network, receivers and servers instead	Section 3, third question; standard 1.2 and the standards' Appendix A
Whether a lost alarm would be noticed	Whether each alarm's delivery is acknowledged or recorded, or only the connection is supervised, what the device and the receiving system show when the connection is lost, and who receives the alert	Section 3, fourth question; standards 4.1 to 4.5
The owner	The owner of the pathway, beside the lead for the network	Section 3, fifth question; standard 7.1
The decision	An independent path, or a shared dependency recorded and accepted, with its interim mitigation and timeline	Section 4, the plan
How it is bought	The contract the pathway falls under, the team that holds it, and the date the contract or its national agreement ends	Section 4, the contracts

Table 3: What a record kept with the self-assessment can hold for each clinical alarm or telemetry pathway.

What a specification can state from the record

For a pathway that needs an independent path, a specification can state, from its record:

1. the pathway, the people who respond and the dependencies it must not share with clinical IT;

2. the time the pathway allows for a switch, shown in two tests that each record whether the switch happened without intervention and whether a test alarm reached the receiving platform: a core failover test, and an end-to-end delivery test for a fault in which the radio stays registered but alarms are not delivered. NHS England's mobile backup guidance notes that an operator account bar can simulate the loss of a network[14], which suits the first test, but an account bar does not reproduce every way an operator core can fail;

3. supplier diversity tested at the operator core as well as at the access network;

4. an acknowledgement or delivery record for each alarm from the receiving platform, and who receives the alert when delivery fails;

5. the dependencies the new path keeps, such as gateway power, backhaul, the private network both paths join and the receiving platform, recorded in the plan;

6. that the new path, as a connection leaving the site, is assessed against the security and compliance standards (8.1 to 8.6) as the site edge is;

7. where one link would serve as both the site's backup and the alarm pathway's upstream path, a link for each role, or the shared use recorded in the plan as a dependency the pathway keeps;

8. the support of the device manufacturer and the receiving platform for carrying the pathway's alarms over the new path, as configured.

5. CSL Group products by site

CSL DualCore® architecture gives a clinical alarm or telemetry pathway two paths through independent operator cores, and CSL supplies both. A deployment specified to deliver the architecture is built to three design principles:

Independent operator cores. The two paths run through mobile operator cores that are independent of each other, rather than through several radio networks that share one core.

Autonomous switching. The SIM, the device or the router makes the switch between cores itself, without waiting for the failed path to recover or for a platform instruction. At scale, reversion to the primary path should be staggered, because displaced devices can otherwise reconnect within a short window and concentrate signalling load on the recovered network.

Private routing on both paths. Private APN routing and an encrypted connection apply on both the primary and the failover path.

For the clinical alarm and telemetry pathways this briefing covers, CSL recommends a further requirement: their traffic does not share the clinical WAN. That requirement follows from what losing those pathways means: a cardiac alarm that does not reach its central station, or a lift alarm call that cannot connect. CSL DualCore architecture can also protect a connection that clinical IT shares with other services, such as an outlying site's backup through a CSL DualSIM Router configured for core independence. Table 4 sets out the upstream path and the CSL Group product for four kinds of device or site. Each product is described in full at www.csl-group.com.

Device or site	Upstream path	CSL product	What it depends on
Medical device with a cellular module	The device's own module, with two multi-network profiles on one rSIM, each on an independent operator core	rSIM	The device manufacturer's approval of the device and modem configuration, and mobile coverage where the device is used
Medical device on the ward LAN	A gateway on the ward network with its own upstream connections	CSL Outpost	The ward LAN between device and gateway, recorded as a dependency the pathway keeps, and the device manufacturer's support for routing its alarms through the gateway to the receiving platform
Small site where full fibre is unavailable or not cost-effective	A link bonded across several mobile operators, with satellite where needed, as the primary connection. When full fibre arrives, it can join the same link, with the mobile connections kept as a backup on a different technology (standard 1.2)	CSL Outpost	Confirmation for each site of what standard 1.1 asks of a mobile or satellite link: performance comparable to full fibre for the intended use, resilience, reliability and latency that meet the site's needs, and a service that can scale with demand; in a new contract, the ability to migrate to full fibre as soon as it becomes viable and cost-effective
Outlying site with one broadband circuit	A cellular second path that does not depend on one operator's core. Cellular failover through one operator gives access diversity only	CSL DualSIM Router, configured for core independence	The router's configuration, and the two operators it uses

Table 4: Upstream paths by device or site, with the CSL product for each and what the choice depends on.

In Table 4, rSIM and the CSL DualSIM Router, in the configuration CSL specifies, implement CSL DualCore architecture. CSL Outpost, in the other two rows, is managed multi-bearer connectivity that complements CSL DualCore architecture. It bonds up to four mobile connections into one link, with satellite and fixed line where needed. The link is independent of any one operator's core only where its connections run through independent operator cores; network names alone do not show this, so it is confirmed for each deployment. A CSL Outpost link can serve as a site's backup connection for clinical IT or as the upstream path for its alarm and telemetry pathways. Where the same link does both, those pathways share an upstream path with clinical IT again, which is the dependency described in section 3. CSL therefore recommends a link for each role. Where a Trust does share one link, including for emergency (break-glass) use of the alarm link, its plan can record the shared use.

CSL's alarm signalling products and services carry fire and security signalling to more than 270 ARCs across Europe, and CSL also provides connectivity services directly to ARCs.

References

[1] NHS England Digital. *Networks and connectivity standards*. Published 13 August 2026; last edited 13 August 2026. Scope, annual self-assessment and the plan where a standard is not met. The count of 21 standards is taken from the eight chapter pages. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards

[2] NHS England Digital. *Networks and connectivity standards: fixed line connectivity*. Last edited 3 August 2026. Standards 1.1 and 1.2, with their DSPT principle mapping. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/fixed-line-connectivity-standards

[3] NHS England Digital. *Networks and connectivity standards: appendices*. Last edited 10 September 2026. Appendix A, supporting technical measures. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/appendices

[4] NHS England Digital. *Networks and connectivity standards: network switching*. Last edited 6 August 2026. Standards 2.1 and 2.2. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/network-switching-standards

[5] NHS England Digital. *Networks and connectivity standards: security and compliance*. Last edited 6 August 2026. Standards 8.1 to 8.6. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/security-and-compliance-standards

[6] NHS England Digital. *Networks and connectivity standards: wireless at NHS sites*. Last edited 6 August 2026. Standards 5.1 to 5.3. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/wireless-at-nhs-sites-standards

[7] NHS England Digital. *Networks and connectivity standards: network cabling*. Last edited 4 August 2026. Standard 3.1. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/network-cabling-standards

[8] NHS England Digital. *Health and Social Care Network*. Last edited 24 August 2026. The page states that many NHS organisations have connected to HSCN. https://digital.nhs.uk/services/health-and-social-care-network

[9] NHS England. *Synnovis ransomware incident* (3 June 2024). Services fully restored by December 2024. Postponement figures from NHS England London's update of 26 September 2024. https://www.england.nhs.uk/synnovis-cyber-incident; https://www.england.nhs.uk/london/2024/09/26/update-on-cyber-incident-clinical-impact-in-south-east-london-thursday-26-september-2024

[10] Portsmouth Hospitals University NHS Trust. *Critical incident declared* (24 June 2026) and *critical incident stood down* (27 June 2026), Queen Alexandra Hospital. https://porthosp.nhs.uk/about-us/news/critical-incident-declared; https://porthosp.nhs.uk/about-us/news/critical-incident-stood-down-portsmouth-hospitals-university-nhs-trust

[11] CSL Group. *The Physics of Critical Connectivity Failure*: the Critical Connectivity Failure Atlas. First edition, October 2026, on Atlas dataset v0.8.16, with incidents coded to 19 September 2026. The corpus is purposive: its percentages describe the selected incidents and do not estimate how often any cause or outcome occurs. The hospital study is Tully, J. L. et al., JAMA Network Open, 2025, as cited in the paper. https://www.csl-group.com/white-papers/the-physics-of-critical-connectivity-failure/; https://doi.org/10.1001/jamanetworkopen.2025.30226

[12] NHS England Digital. *Networks and connectivity centre of excellence*. Last edited 10 September 2026. The page states that the Networks and Connectivity Programme "is part of Frontline Productivity portfolio of work", and the standards page that they "have been produced by the Frontline Productivity Networks and Connectivity programme in NHS England". https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence

[13] NHS England Digital. *Networks and connectivity standards: leadership and governance*. Last edited 6 August 2026. Standard 7.1. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/networks-and-connectivity-standards/leadership-and-governance-standards

[14] NHS England Digital. *Mobile backup solutions for fixed healthcare sites*: potential technologies and solutions, and business continuity. Both last edited 19 March 2025. https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/knowledge-products-and-artefacts/advice-and-guidance/mobile-backup-solutions-for-fixed-healthcare-sites/potential-technologies-and-solutions; https://digital.nhs.uk/services/networks-and-connectivity-centre-of-excellence/knowledge-products-and-artefacts/advice-and-guidance/mobile-backup-solutions-for-fixed-healthcare-sites/business-continuity

[15] Government Commercial Agency (GCA). Procurement routes: HSCN DPS RM3825, to 23 May 2028; RM6261 Mobile Voice and Data Services, to 12 December 2026; RM6391 Mobile Voice and Data Services 2, from 17 September 2026 to 16 September 2028; RM6116 Network Services 3, to 17 July 2027; RM6377 Network Services 4, which GCA says "will replace RM6116 Network Services 3", framework award estimated for March/April 2027, with product types that include IoT and connected places (including security and CCTV), paging and alerting services, and critical communications. Crown Commercial Service became the GCA on 1 April 2026. https://www.gca.gov.uk/agreements/RM3825; https://www.gca.gov.uk/agreements/RM6261; https://contractsfinder.service.gov.uk/Notice/b276ceba-9e91-4984-b56e-796829c83337; https://www.gca.gov.uk/agreements/RM6391; https://www.gca.gov.uk/agreements/RM6116; https://www.gca.gov.uk/agreements/RM6377

[16] DSIT. *Telecommunications modernisation: connectivity timeline*. Published 24 March 2026; last updated 29 July 2026. https://gov.uk/guidance/telecommunications-modernisation-connectivity-timeline

[17] DSIT. *Moving landlines to digital technologies*. Published 28 May 2025. https://gov.uk/guidance/moving-landlines-to-digital-technologies

[18] DSIT. *Telecoms modernisation: critical national infrastructure charter*. Published 24 March 2026; last updated 13 October 2026. https://gov.uk/government/publications/telecoms-modernisation-critical-national-infrastructure-charter/telecoms-modernisation-critical-national-infrastructure-charter

[19] DSIT. *Preparing for the 2G switch-off: mobile phones and smart devices*. Published 24 March 2026; last updated 29 July 2026. https://gov.uk/guidance/preparing-for-the-2g-switch-off-mobile-phones-and-smart-devices

[20] NHS England Digital. *Guidance on protecting connected medical devices: Step 1. Identify connected medical devices*. Last edited 13 November 2025. https://digital.nhs.uk/cyber-and-data-security/guidance-and-resources/guidance-on-protecting-connected-medical-devices/step-1.-identify-connected-medical-devices

CSL

Healthcare Connectivity Architecture

Secondary Care | Briefing

General enquiries

To discuss the clinical alarm and telemetry pathways on a Trust's estate, contact sales@csl-group.com.

csl-group.com

+44 (0) 1895 474 474

This publication

Briefing, October 2026. Evidence cut-off: 22 October 2026; CSL product information confirmed to 28 October 2026.

Disclaimer This document is for general informational purposes only and does not constitute professional, medical, regulatory, financial, or legal advice. Service-level commitments and SLA terms are contract-specific.

© 2026 CSL Group. All rights reserved. rSIM, CSL DualCore, CSL DualSIM and Critical Connectivity are registered trademarks of CSL (DualCom) Limited. The CSL mark is a trademark of CSL Group. Other names are the property of their respective owners and are used for identification only.