Reliability of a cooling service from enumerated outage states scored by heat-flow physics, on a two-hall data center with a shared electrical feeder

Ricardo Marquez
rmarquez123@gmail.com

November 10, 2026

Abstract

A pair of data-hall cooling units is declared redundant when each could carry the other hall's load, and the claim becomes a number through a reliability block diagram in which the units are drawn in parallel or in series. A wrong structure is silent: the claim reads the same on a dashboard whether the units back each other up or not, and the first evidence against it is the day a unit fails and a hall overheats. This report is a forward study of that step on a two-hall system with one wall opening between the halls and one electrical feeder: nothing is estimated from data, and every result follows from stated inputs. Two processes are run on the same component list. In the first the analyst asserts the success structure, which combinations of failed components lose the service, and the arithmetic follows exactly. In the second every combination is enumerated with its probability and scored by the heat-flow physics of the halls, so that the structure is computed rather than asserted. The two admissible assertions of the first process differ by a factor of about 39 in service unavailability, and nothing in the component list says which is right. The second process reproduces each assertion exactly, each at its own wall opening: the units are in series when the opening cannot pass one hall's load and in parallel when it can, and the switch happens at that load. Widening the boundary to include the feeder recovers the product of supply and equipment reliabilities as a result rather than an input, and a common-cause failure, one disturbance that fails supply and equipment together, makes that product overstate the service reliability. The minimal cut sets and path sets of the computed structure give the fault tree and the block diagram as outputs, and the tree changes shape at the same load. The redundancy of a pair of units is a property of the network between them, and the opening's capacity is the quantity to verify before the claim is made. The system has two to three components, every state is enumerated, and the computation is exact.

1 Introduction

Data centers promise their tenants an availability, and the promise rests on redundancy: for every piece of equipment whose loss would interrupt the service there is a spare, written as $N + 1$. The cooling system is where the promise is most exposed, because a hall that loses its cooling overheats within minutes, and because the redundancy of cooling equipment is harder to verify than that of electrical equipment. Two chillers on one bus either carry the load or they do not; two cooling units in two halls can each carry both loads on paper and still fail to back each other up, because the heat of one hall has to reach the unit of the other through whatever opening connects them. Availability claims therefore depend on how the equipment list is turned into a model of the service, a step that is routine, rarely audited, and decisive.

The standard tools for that step are the reliability block diagram and the fault tree. Both are exact once the success structure of the system is given: the block diagram draws the components in series and parallel and multiplies and complements their reliabilities, the fault tree writes the top event as a logic of component failures and evaluates it through its minimal cut sets (Rausand and Høyland 2004; Billinton and Allan 1992; Vesely et al. 1981). The structure itself is supplied by the analyst from the equipment list and a systematic analysis of its failure modes, and the arithmetic that follows cannot reveal a wrong structure, because it is correct for whichever structure it is given. Where the service depends on flows through a network rather than on the presence of components, the network-reliability literature replaces the asserted structure by a connectivity or flow condition evaluated on each state of the network (Ball et al. 1995), and the adequacy assessment of power systems scores every outage state by a load-flow solution and a load-shedding rule rather than by a diagram (Billinton and Allan 1992). Industrial practice for continuous-power facilities retains the minimal cut-set method, the series and parallel formulas of its recommended practice, with the block diagram and the fault tree as its drawings (IEEE 2013; IEEE 2025). When components share a cause of failure, the independence that the diagram's products assume fails as well, and the common-cause literature treats the common-cause failure as an explicit part of the model (Mosleh et al. 1988).

TERRA is a general-purpose engine for probabilistic inference on physical networks. It treats a network as a graph on which the physics is written locally: a conservation law at every node, a constitutive law on every edge, and the boundary and component laws that close them, attached to the same graph. Each of these is a relation the state of the network must satisfy, and the engine assembles them and solves them together. Each physical domain brings its own laws, units and file formats; the solving and the bookkeeping are shared. This study builds a two-hall thermal network with it, scores every outage state of the network's components by the engine's steady heat-flow solution, and sums the probabilities of the states that still deliver the cooling service. Section 3 states the framework and each of these constructs precisely.

The quantity throughout is the service reliability, R_{service}: the probability that the cooling service is delivered at a random hour, an availability in the sense of IEEE Std 3006.7 (IEEE 2013) rather than a survival probability over a mission time; the report keeps the name under which the claim is made. This report performs the step from the equipment list to the service reliability twice, on the smallest system on which the two processes can disagree, and compares the processes rather than the numbers. The reason to compare them is that this step is where a redundancy claim is decided and where nothing checks it: an asserted structure costs nothing until it costs an outage, and the second process makes the structure a computed property of the network, which a measurement can contradict and a drawing cannot. The question is whether two cooling units that could each carry both halls form a redundant pair, and what the reliability of the cooling service is. The construction is a two-hall thermal network, a two-state model of each component, and the enumeration of every outage state with its heat-flow solution: the practice of power-system adequacy assessment, where a load-flow solution scores each state, applied to a thermal network, on which the same enumeration and the same scoring run unchanged. Nothing in the construction is new. What the report adds is a demonstration that the success structure a block diagram asks for is a computable property of the network, that the study boundary is the list of components enumerated, and that the product rule for supply and equipment is a consequence with stated conditions. Section 2 describes the system and its inputs, Section 3 the two processes and the relations they rest on, Section 4 the results, Section 5 their limits, and Section 6 the answers; the appendices carry the reproduction recipe, the data statement and the notation.

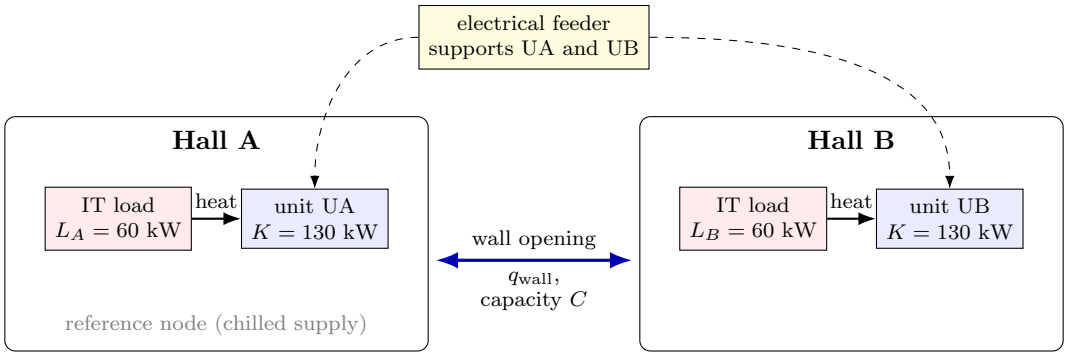


Figure 1: The system. Two data halls, each with an IT heat load of 60 kW and one cooling unit of 130 kW capacity, connected by a wall opening through which heat crosses when the halls' temperatures differ, up to the opening's heat-flow capacity C. One electrical feeder supports both units (dashed arrows): its outage takes both out of service. Hall A's chilled supply is the reference node of the thermal network.

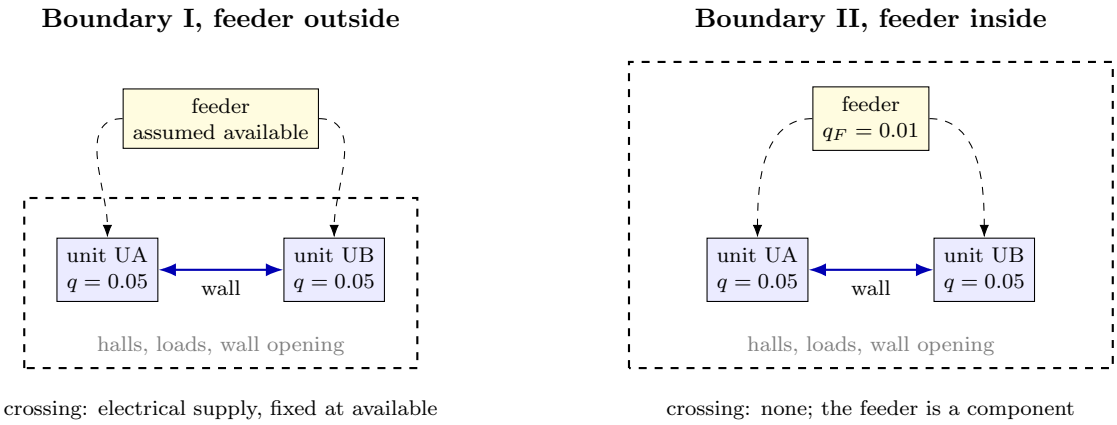


Figure 2: The two study boundaries drawn on the same system (dashed rectangles). Under boundary I the electrical supply crosses the boundary and is held fixed at available, which is the convention of a cooling-equipment study; under boundary II the feeder is a component with its own unavailability and nothing crosses. The boundary is the list of components the enumeration ranges over.

2 System and data

The system. Figure 1 shows it. Two data halls, A and B, each carry an IT heat load of 60 kW and each has one cooling unit of 130 kW capacity, running at 60 kW of duty at the operating point, so that either unit alone has the capacity for both halls. A wall opening connects the halls: heat crosses it in proportion to the temperature difference between the halls, and the opening carries at most its heat-flow capacity C, the one parameter the study varies. Hall A's chilled supply is the reference node of the thermal network. One electrical feeder supports both units, and its outage takes both out of service. The industry's recommended practice for reliability planning (IEEE 2025) stops at the 480 V point of use, with failure defined as a loss of incoming power for more than 5 s; the cooling service lies outside that boundary, and this study begins where the practice ends.

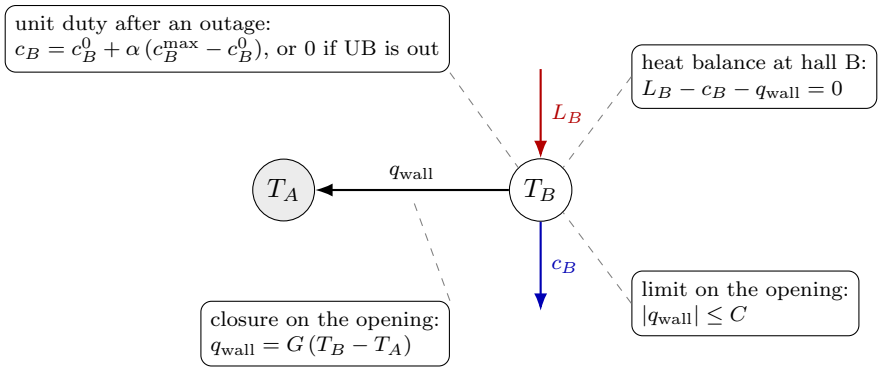


Figure 3: The model around one hall. The unknowns are the hall temperatures and the heat flow on the opening; the relations are the heat balance at each hall, the closure law on the opening, the duty of each unit after an outage (its own duty plus a share α of its headroom, or zero when it is out), and the capacity limit that decides whether the state delivers the service. There are no readings and nothing is estimated.

Table 1: Design inputs of the study. Unavailabilities are hour-level: the probability that a component is out of service at a random hour. Every result in the report is computed from these values.

quantity	symbol	unit	value
IT heat load per hall	L_i	kW	60
cooling unit capacity, each	$c_i^{\max}$	kW	130
cooling unit duty at the operating point	c_i^0	kW	60
allowable hall temperature rise	ΔT_{allow}	K	8
wall opening conductance, narrow / wide	G	W/K	3750 / 10000
wall opening capacity, narrow / wide	C	kW	30 / 80
wall opening capacity sweep	C	kW	20 to 100 in steps of 10
unavailability of each cooling unit	q	–	0.05
unavailability of the feeder	q_F	–	0.01
common-cause state, calm, weight	π_{calm}	–	0.9
common-cause state, disturbance, weight	π_{dist}	–	0.1
disturbance-state unavailabilities, feeder / each unit	$q_F(w)$, $q_i(w)$	–	0.08 / 0.30

The boundary. The two boundaries the study uses are drawn on the same system in Figure 2. Under boundary I the feeder lies outside: the electrical supply crosses the boundary and is held fixed at available, which is the convention of a study of the cooling equipment alone. Under boundary II the feeder is a component with its own unavailability and nothing crosses. The boundary is not a drawing convention here; it is the list of components whose outage states are enumerated, and Section 3.6 states what moving the feeder across it does to the result.

The model. The relations around one hall are shown in Figure 3. The unknowns are the two hall temperatures and the heat flow on the opening. The relations are a heat balance at each hall, the closure law on the opening, the duty of each unit after an outage, and the capacity limit on the opening that decides whether a state delivers the service. There are no readings, no priors and nothing estimated: the study is a forward computation on stated inputs.

The data. Every input is listed in Table 1. The study is a constructed design case with no telemetry; the unavailabilities are design numbers of the kind a reliability data handbook supplies (IEEE Std 3006.8, IEEE 2018), hour-level, meaning the probability that the component is out of

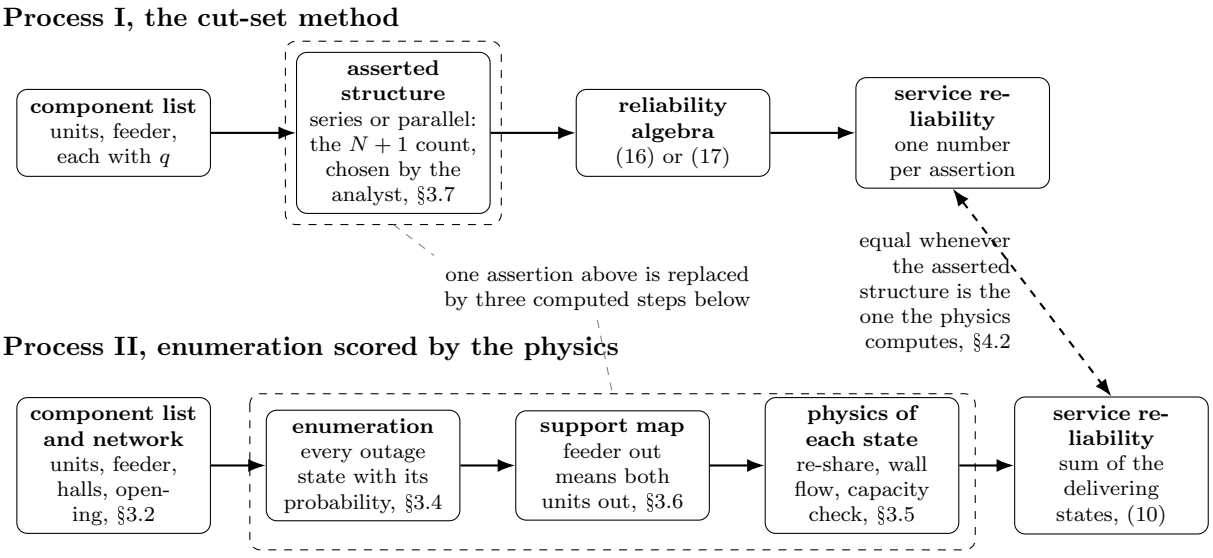


Figure 4: The two processes that produce a service reliability from the same component list. Process I asks the analyst to assert the success structure (which combinations of failed components lose the service), which is where an $N + 1$ count enters, and then applies exact algebra to it. Process II enumerates every combination, applies the feeder's support map, scores each combination with the heat-flow physics, and sums the probabilities of the combinations that deliver. The one asserted step of the first process is replaced by the three computed steps of the second (dashed frames), and the two outputs coincide whenever the assertion was right (dashed double arrow): the second process contains the first and checks its assertion.

service at a random hour. The assumptions are these. The thermal physics is steady and linear, with no thermal mass and no time. The limit checked on the opening is a heat-flow capacity, not a hall temperature limit. A unit is either fully available or fully out; a failed unit's duty is taken up by the surviving units in proportion to their headroom, the convention the engine uses for generators in a power network, and any duty beyond the total headroom is left unserved. The feeder's effect on the units is declared as a dependency, both units out, not solved as an electrical network. Components fail independently except in the common-cause-failure variant, which is a mixture of two independent models.

3 Methods

3.1 Overview

Figure 4 lays out the two processes that turn the component list of Table 1 into a service reliability. Process I is the cut-set method, the block diagram read as its cut sets: the analyst decides whether the two units are in series or in parallel, and the reliability algebra of Section 3.7 gives one number per decision. Process II enumerates every outage state of the components (Section 3.4), applies the feeder's support map (Section 3.6), scores each state by the heat-flow physics of Section 3.2 through the margin of Section 3.5, and sums the probabilities of the states that deliver. The two processes share their first and last boxes and differ in one step: in the first the success structure is asserted, in the second it is computed. Section 3.8 reads the fault tree and the block diagram off the computed structure, so that the artifacts the first process starts from are outputs of the second. Section 3.9 extends the second process to a common-cause failure, and Section 3.10 defines the measures the results are read by.

3.2 The heat flow between the halls

The halls are the nodes of a steady linear thermal network and the opening is its one edge. At each hall the heat load must leave either through the hall's own unit or through the opening, so with the flow on the opening counted from B to A, the heat balance at hall B is

$$L_B - c_B - q_{\text{wall}} = 0, \tag{1}$$

where L_B is hall B's IT load, c_B the duty of its unit and q_{wall} the heat crossing the opening, all in watts; hall A's balance has the same form with the sign of q_{wall} reversed. The flow on the opening follows the temperature difference between the halls through the opening's conductance,

$$q_{\text{wall}} = G\,(T_B - T_A), \tag{2}$$

with G in W/K and T_A, T_B the hall temperatures. This is the closure law of the edge, the thermal counterpart of a conductance in a circuit, and it assumes a steady state with no thermal mass. The opening carries at most its heat-flow capacity,

$$|q_{\text{wall}}| \le C, \tag{3}$$

a limit on the edge rather than a relation the state must satisfy: a state in which the flow exceeds C is a state in which the service fails, and C is the one input the study sweeps. The conductance is not a separate input: it is the capacity over the hall temperature rise the IT equipment tolerates,

$$G = C/\Delta T_{\text{allow}}, \tag{4}$$

with $\Delta T_{\text{allow}} = 8$ K here, so that the capacity check (3) and the statement that the halls stay within the allowable rise are one statement; Table 1 lists the conductance of each opening. After an outage the surviving units take up the lost duty in proportion to their headroom. With L the duty the failed units dropped and H the headroom of the surviving units, both in watts, each survivor runs at its operating-point duty plus the share

$$\alpha = \frac{\min(L, H)}{H} \tag{5}$$

of its own headroom. With $c_i^{\max}$ the capacity of the unit at hall i and c_i^0 its duty at the operating point, its duty after the outage is

$$c_i = c_i^0 + \alpha\,(c_i^{\max} - c_i^0), \tag{6}$$

or zero when the unit itself is out; this is the relation drawn beside hall B in Figure 3. The headroom H is the sum of $c_i^{\max} - c_i^0$ over the surviving units and the lost duty L the sum of c_i^0 over the failed ones, so with every survivor at (6) the lost duty is taken up in full whenever it fits. The duty that no headroom can absorb is the shortfall

$$D = \max(L - H, 0). \tag{7}$$

The re-share rule is the convention the engine uses for generators in a power network, adopted here unchanged; it says nothing about how fast the survivors respond, which a steady model cannot. With one unit out, $L = 60$ kW and $H = 70$ kW, so α is below one and $D = 0$ by (7): the surviving unit can carry both halls, and by (1) the whole of the lost duty, 60 kW, must cross the opening. With both units out, $H = 0$ and $D = 120$ kW.

3.3 The network model

These relations are an instance of a general construction. A physical network is a graph, and the engine writes its physics on that graph. The state is a vector z of unknowns: one potential per node, such as a temperature or a head, one flow per edge, and whatever else the physics determines. Every relation the state must satisfy — a conservation law at a node, a constitutive law on an edge, a boundary condition, a component law — is written as a residual $r_k(z)$ that vanishes when the relation holds. With as many independent relations as unknowns and no readings to reconcile, the state is the unique solution of

$$r(z) = 0, \tag{8}$$

found by Newton iteration on the sparse Jacobian J of the residuals, $J_{ki} = \partial r_k / \partial z_i$, from a stated start until the residual norm falls below a tolerance ϵ:

$$z \leftarrow z - J^{-1} r(z) \quad \text{until } \|r(z)\| < \epsilon. \tag{9}$$

A transient model adds a storage term to the conservation relations and advances the state in time with an implicit integrator, solving a system of the same kind at every step. Nothing in this report is estimated from data: every input is stated, and the results are the solutions of these relations and quantities computed from them. Everything that follows is a choice of unknowns and relations on this graph.

3.4 The service reliability as a sum over outage states

The quantity the second process owes is the service reliability, the probability that the cooling service is delivered at a random hour. It is the total probability of the outage states in which the service is delivered,

$$R_{\text{service}} = \sum_{S} P(S)\, \mathbf{1}[S], \tag{10}$$

with S an outage state, the set of components out of service at once, $P(S)$ its probability and $\mathbf{1}[S]$ its verdict, one when the service is delivered in that state and zero when it is lost. The derivation runs from the sum down to its two factors: the probability of a state here, the verdict in Section 3.5. Figure 5 follows one pass of it per listed state.

Each component is in one of two states, available or out of service, and is out with its unavailability q, the probability that it is down at a random hour. When the components fail independently, the probability of S is the product of the unavailabilities of the components in S and the availabilities of the rest,

$$P(S) = \prod_{i \in S} q_i \prod_{i \notin S} (1 - q_i). \tag{11}$$

Here $\prod_{i \in S}$ multiplies the terms that follow it over every component i in the set S, and $\prod_{i \notin S}$ over every component not in S; with two units of unavailability q and nothing else, the state in which both are out has probability q^2 and the state in which neither is out has $(1 - q)^2$. Writing $P(S)$ as a product is the independence assumption of the block diagram made explicit, and Section 3.9 relaxes it. The enumeration lists the states in order of decreasing probability and may stop before the list is complete; the probability mass of the states it has not listed,

$$\varepsilon = 1 - \sum_{S \text{ listed}} P(S), \tag{12}$$

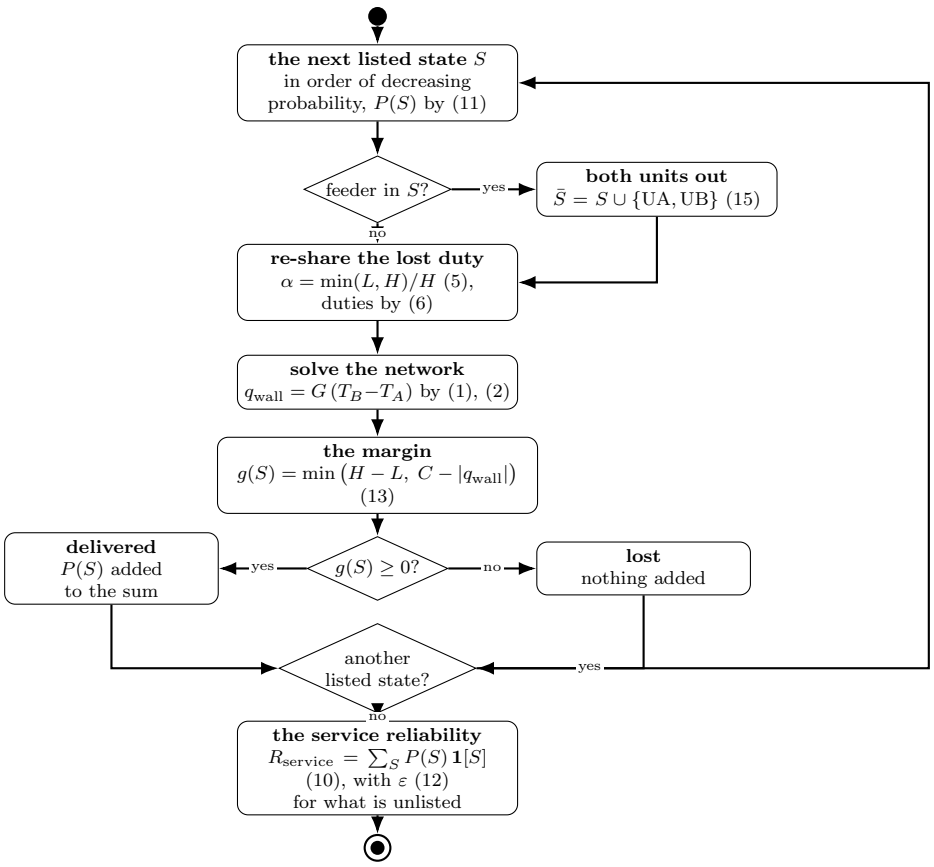


Figure 5: The steps from the component list to the service reliability, one pass per outage state. The diagram runs from the filled start node to the ringed end node. Each listed state is taken in order of decreasing probability; when the feeder is in it, both units are added to the outage by the feeder dependency (15); the lost duty is re-shared by (5), the network is solved for the flow on the opening, and the margin (13) decides whether the state's probability joins the sum (10). The loop closes over the list, and the mass of the states never listed, (12), bounds what the sum misses. The two decisions are the two places the physics, not the analyst, decides.

bounds whatever those states could contribute to any probability computed from the list, because every such probability is a sum of non-negative terms; on a truncated list (10) is a lower bound and ε bounds what it misses. Here the list is complete, 4 states under boundary I and 8 under boundary II, so ε is round-off.

3.5 The margin and the verdict of a state

The verdict of (10) is computed from the physics of the state. For each listed state the components in S are removed, the lost duty is re-shared by (5), and the network of Section 3.2 is solved for the flow on the opening. Two things can then fail the service, duty left unserved and heat on the opening beyond its capacity, and the state's margin is the smaller of its two distances to those failures,

$$g(S) = \min\big(H(S) - L(S),\ C - |q_{\mathrm{wall}}(S)|\big), \tag{13}$$

in watts, with $H(S) - L(S)$ the headroom the survivors have left once the lost duty is re-shared, which is $-D(S)$ of (7) when that duty does not fit, and $q_{\mathrm{wall}}(S)$ the flow (2) in that state. The

service is delivered in state S when the margin is non-negative, which the indicator

$$\mathbf{1}[S] = \begin{cases} 1 & g(S) \geq 0 \\ 0 & g(S) < 0 \end{cases} \tag{14}$$

records as one or zero; an indicator is a quantity that is one when its condition holds and zero otherwise, and the brace lists the two cases. This is the step at which the two processes part. The block diagram asks the analyst for the set of states in which the service survives; (14) computes it from the physics, and a change to the opening's capacity changes the set without anything being redrawn.

3.6 The boundary and the feeder dependency

The boundary of the study is the list of components over which (11) ranges. A component left off the list is held available in every state, which is what a cooling-equipment study does with the electrical supply. Under boundary II the feeder is on the list, and its effect on the units enters through a support map: in any state that contains the feeder, both units are added to the set of units out of service, so the state the re-share (5) and the margin (13) see is

$$\bar{S} = \begin{cases} S \cup \{\mathrm{UA}, \mathrm{UB}\} & \text{feeder} \in S \\ S & \text{feeder} \notin S, \end{cases} \tag{15}$$

with UA and UB the two units; the probability of the state stays $P(S)$ of (11), since the units did not fail, they lost their supply. The map declares the dependency; it does not solve an electrical network, and Section 5 returns to that limit.

3.7 The series and parallel formulas

Process I is stated in the same terms, the series and parallel formulas of the cut-set method (IEEE Std 3006.5, IEEE 2014). For two units of unavailability q, the series assertion holds that both units are needed, so the service survives only when neither is out,

$$R_{\mathrm{series}} = (1 - q)^2, \tag{16}$$

while the parallel assertion holds that either unit suffices, so the service is lost only when both are out,

$$R_{\mathrm{parallel}} = 1 - q^2. \tag{17}$$

When a feeder of unavailability q_F that both units need is added in series with the pair, the diagram multiplies the feeder's reliability $R_F = 1 - q_F$ by the pair's reliability R_C, whichever assertion produced R_C,

$$R_{\mathrm{service}} = R_F \, R_C. \tag{18}$$

Each of the three is exact for the structure it assumes, and the product also assumes that the feeder and the units fail independently. The report's question is not whether the algebra is right but which structure the system has, and whether that independence holds.

3.8 The cut sets, path sets and fault tree from the enumeration

The rule that assigns every outage state its verdict, delivered or lost, is called the structure function of the system; it is the object a block diagram or a fault tree encodes, and here it is the indicator (14) computed from the physics. The artifacts of the first process follow from it, and the notation needed to say so is small: a set of components is written with a capital letter, K or P; U is the full component list; $K' \subsetneq K$ reads "K' is a proper subset of K", a set made of some but not all of K's components; $S \subseteq S'$ reads "S is contained in S'"; and $U \setminus P$ is the list with the components of P removed, that is, the outage state in which everything but P is out. A minimal cut set K is a set of components whose joint loss loses the service while the loss of any smaller part of K does not: its own verdict is zero and the verdict of every proper subset is one,

$$\mathbf{1}[K] < \min_{K' \subsetneq K} \mathbf{1}[K'], \tag{19}$$

where min over K' is the smallest verdict among the proper subsets, so the inequality says that all of them deliver. Every component in K is therefore needed for the loss. A minimal path set P is the mirror image: a set of components whose availability alone keeps the service, while no smaller part of P does. With everything but P out the service is delivered, and with everything but a proper part of P out it is lost,

$$\mathbf{1}[U \setminus P] > \max_{P' \subsetneq P} \mathbf{1}[U \setminus P']. \tag{20}$$

where max over P' is the largest verdict among the proper subsets, so the inequality says that none of them delivers. Cut sets and path sets each describe the structure completely, and they are unique, when the structure is coherent, which means that taking out more components can never turn a lost service back into a delivered one: whenever one state S is contained in a larger state S', the verdict of S is at least that of S',

$$\mathbf{1}[S] \geq \mathbf{1}[S'] \quad \text{whenever } S \subseteq S'. \tag{21}$$

A drawn diagram assumes coherence; a structure computed from physics can violate it, so the condition is checked on every listed state and a violation is reported with the pair of states that witness it. With the cut sets in hand, the fault tree follows. Its top event is the loss of the service, and it is lost exactly when, for at least one minimal cut set, every component in that set is out: in the notation of logic, an OR ($\bigvee$) over the cut sets of an AND ($\bigwedge$) over the components of each (Vesely et al. 1981),

$$\text{service lost} = \bigvee_{K \in \mathcal{K}} \bigwedge_{i \in K} \{\text{component } i \text{ out}\}, \tag{22}$$

with $\mathcal{K}$ the collection of minimal cut sets. Each AND is one gate of the tree, and the block diagram is the dual picture, one parallel branch per minimal path set. The derivation uses the enumeration and nothing else. Every proper subset of a listed state is more probable than the state and is listed before it, so a listing truncated at the mass ε of (12) misses no minimal cut set more probable than the least probable state it lists, and the probability of all the cut sets it misses together is at most ε; the listing here is complete and misses none. Path sets are read from the largest delivered states, which a truncation drops first, and are derived from complete listings only. The tree is derived at both openings and at every capacity of the sweep, so that the capacity at which it changes shape is read off the trees themselves. This is the pattern of probabilistic risk assessment in the nuclear industry, where the success criteria of a safety function are fixed by a physics calculation before the trees are drawn (USNRC 1983), with the criteria computed here for every state rather than once.

3.9 Common-cause failure

The product (18) rests on the feeder and the units failing independently. A disturbance that both interrupts the supply and trips the units on restart breaks that: it is a common-cause failure, a single shared cause that fails several components at once. The study represents such a cause as a discrete state w with weight π_w; in each state of the cause the components are independent with their own unavailabilities $q_i(w)$, so the probability of an outage state is a mixture of product models, that is, a weighted average of the product (11) evaluated under each state of the cause, with the weights π_w summing to one,

$$P(S) = \sum_{w} \pi_w \prod_{i \in S} q_i(w) \prod_{i \notin S} \big(1 - q_i(w)\big), \tag{23}$$

where $\sum_w$ adds the terms over the states of the cause. The service reliability under the cause is the weighted sum of the two enumerations, and because the physics of a state depends on S alone, both enumerations share every solved state. The cooling-equipment reliability conditional on the feeder being up is the service reliability divided by the feeder's,

$$R_{C|F} = \frac{R_{\text{service}}}{R_F}, \tag{24}$$

which equals the unconditional R_C only when the cause couples nothing, so the gap between the two measures the coupling. The calm state carries weight 0.9 and the unavailabilities of Table 1; the disturbance state carries weight 0.1 and raises the feeder to 0.08 and each unit to 0.30.

3.10 Metrics

Four measures read the results. The service reliability (10) answers whether the service is delivered, and is read against the two assertions (16) and (17). How far an $N+1$ claim made on the parallel assertion understates the risk when the series structure is the truth is measured by the ratio of the unavailability on the narrow opening to that on the wide one,

$$\rho = \frac{1 - R_{\text{service}}^{\text{narrow}}}{1 - R_{\text{service}}^{\text{wide}}}, \tag{25}$$

whose reference is one, the value at which the claim is right. An hour-level unavailability is translated into an exposure an operator can read, in hours per year,

$$H = 8760\,(1 - R_{\text{service}}), \tag{26}$$

read against the claim's own figure. And an enumerated reliability is compared with the closed form of the structure it is claimed to reproduce, that is, the block-diagram formula evaluated by hand, by their absolute difference,

$$\delta = \big|R_{\text{service}} - R_{\text{closed form}}\big|, \tag{27}$$

which is reported as a bound, because its value is round-off.

4 Results

4.1 The two block-diagram assertions

Table 2 sets out the two structures an analyst can assert for the pair and the service reliability each gives by (16) and (17). The equipment list supports both: each unit has 130 kW against a combined

Table 2: The two admissible block-diagram assertions for two units of unavailability $q = 0.05$, and the service reliability each yields. The unavailabilities differ by the factor $\rho = 39$ of (25).

asserted structure	relation	R_{service} [–]	$1 - R_{\text{service}}$ [–]
parallel, either unit suffices	(17)	0.9975	0.0025
series, both units needed	(16)	0.9025	0.0975

Table 3: The four outage states under boundary I, their probabilities from (11), the heat the opening must carry and the duty left unserved in each, and the verdict of (14) on each opening. Service reliability by (10): 0.9025 on the narrow opening and 0.9975 on the wide one; un-enumerated mass below 10^{-12} in both.

outage state	$P(S)$ [–]	wall flow [kW]	shortfall [kW]	narrow (30 kW)	wide (80 kW)
none	0.9025	0	0	delivered	delivered
unit A	0.0475	60	0	lost	delivered
unit B	0.0475	60	0	lost	delivered
both units	0.0025	60	120	lost	lost

load of 120 kW. The unavailabilities differ by the factor $\rho = 39$ of (25), against a reference of one. Nothing in the first process distinguishes the two; the choice is made when the diagram is drawn, and the arithmetic that follows is exact for either. For the owner, the gap between the two rows is the risk an unverified count carries.

4.2 The enumeration at two openings

The 4 outage states under boundary I are listed in Table 3 with their probabilities from (11), the heat the opening must carry in each, the shortfall, and the verdict of (14) on each opening. With one unit out, the other hall's 60 kW must cross the opening and the shortfall is zero; with both out the shortfall is 120 kW on either opening. On the narrow opening of 30 kW both single-unit states are lost, because 60 kW exceeds the capacity, and the service reliability by (10) is 0.9025, which equals the series assertion (16) to below 10^{-12} by (27). On the wide opening of 80 kW both single-unit states are delivered and the service reliability is 0.9975, which equals the parallel assertion (17) to below 10^{-12}. The same two units and the same unavailabilities give the series answer at one opening and the parallel answer at the other; the structure that the first process asks the analyst to assert is here read off the physics, and both assertions turn out to be right, each for its own opening. Figure 6 plots the service reliability of (10) against the opening's capacity swept from 20 to 100 kW, with the two assertions as horizontal lines and one hall's load as the vertical line. The result is a single step: the physics sits on the series line up to 50 kW and on the parallel line from 60 kW, and no intermediate value occurs, because the only state whose verdict depends on the capacity is the single-unit outage, and its verdict flips when the capacity reaches the 60 kW that must cross. The step is at one hall's load, which is the number a designer can act on. The exposure of (26) is about 22 hours per year on the parallel side of the step and about 854 on the series side. For the owner, the opening's capacity is the quantity to verify before the pair is called redundant, and the sweep names the value it must exceed.

4.3 The feeder inside the boundary

The 8 outage states under boundary II are listed in Table 4, with the support map applied in its third column, and Figure 7 plots the same states by probability with their verdicts on both openings.

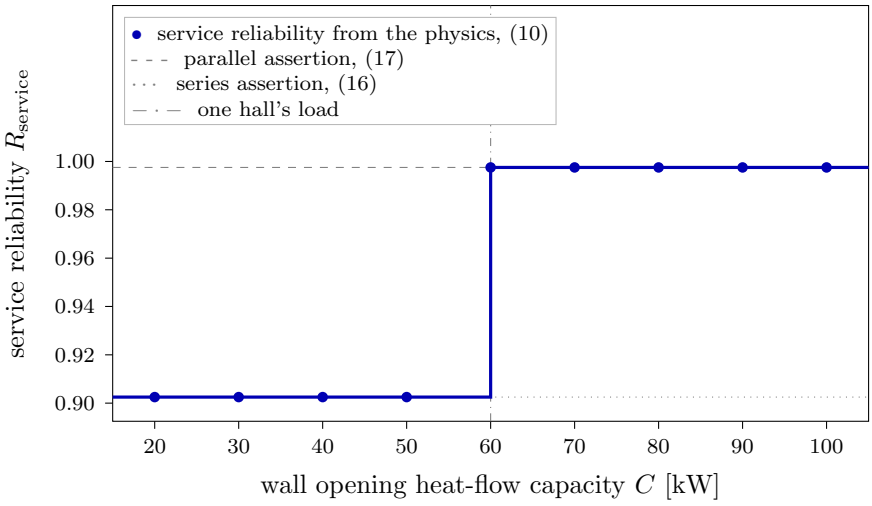


Figure 6: Service reliability against the heat-flow capacity of the wall opening, swept from 20 to 100 kW in steps of 10 kW (dots, joined as a step). The gray lines are the two block-diagram assertions and the vertical line is one hall's load. The physics sits on the series line up to 50 kW and on the parallel line from 60 kW; nothing in between occurs. The riser is drawn at 60 kW, where the limit (3) first admits one hall's load with a unit out; the sweep's grid alone would only place it between 50 and 60 kW.

Table 4: The eight outage states under boundary II. The third column applies the support map: a feeder outage puts both units out of service. Service reliability by (10): 0.8935 on the narrow opening and 0.9875 on the wide one; un-enumerated mass below 10^{-12} in both.

outage state	$P(S)$ [–]	units out of service	narrow	wide
none	0.8935	none	delivered	delivered
unit A	0.0470	UA	lost	delivered
unit B	0.0470	UB	lost	delivered
both units	0.0025	UA, UB	lost	lost
feeder	0.0090	UA, UB	lost	lost
unit A, feeder	0.0005	UA, UB	lost	lost
unit B, feeder	0.0005	UA, UB	lost	lost
both units, feeder	0.0000	UA, UB	lost	lost

Every state containing the feeder is lost on both openings, because the map puts both units out and the shortfall is 120 kW; the two single-unit states are the only ones whose verdict depends on the opening, as under boundary I. The service reliability by (10) is 0.9875 on the wide opening and 0.8935 on the narrow one. Each equals the product (18) of the feeder's reliability $R_F = 0.99$ and the pair's reliability under boundary I, 0.9975 and 0.9025, to below 10^{-12} and below 10^{-12} by (27). The product rule of the block diagram is recovered as a result of the enumeration, with its two conditions visible: the feeder is on the component list, and the components are independent. A study under boundary I reports 0.9975, which is the cooling-equipment reliability conditional on power being available, a valid number with a stated assumption, and not the reliability of the delivered service, which is smaller by the feeder's factor. For the owner, leaving the feeder out of the model is a choice the model makes visible, and the product form says what that choice costs.

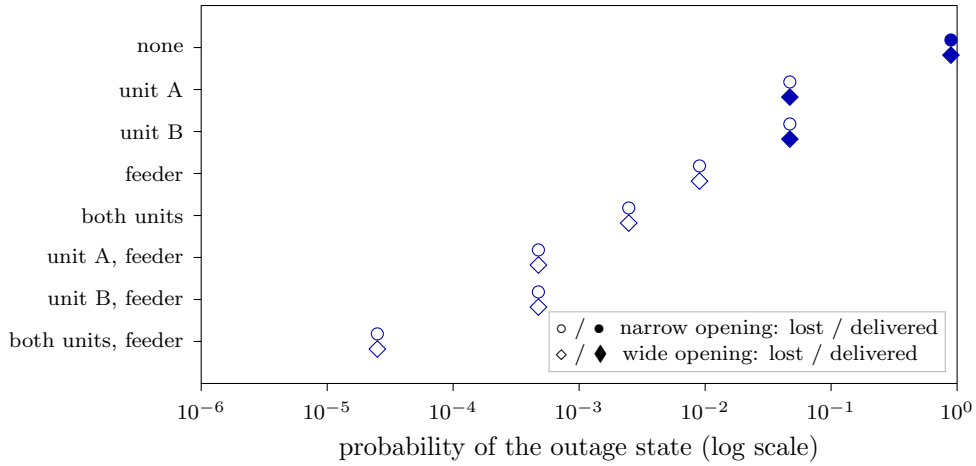


Figure 7: The eight outage states of boundary II, ordered by probability (the smallest, all three components out, is drawn at the axis floor). For each state the upper mark is the verdict on the narrow opening and the lower mark the verdict on the wide one; a filled mark means the service is delivered. Only the two single-unit states change verdict between the openings; every state with the feeder out is lost on both.

Table 5: The success structure read off the enumeration under boundary II at the two openings: the minimal cut sets of (19), the minimal path sets of (20), the coherence check (21), and the bound on the probability mass of any cut set the listing could miss, which is round-off because the listing is complete. The cut sets are the fault tree of Figure 8 and the path sets are the branches of the block diagram.

opening [kW]	minimal cut sets	minimal path sets	coherent	mass not listed [–]
30	{UA}, {UB}, {feeder}	{UA, UB, feeder}	yes	below 10^{-12}
80	{feeder}, {UA, UB}	{UA, feeder}, {UB, feeder}	yes	below 10^{-12}

4.4 The derived fault tree

The minimal cut sets and minimal path sets read off the enumeration of Table 4 by (19) and (20) are collected in Table 5, and Figure 8 shows the fault tree of each by (22). The structure satisfies (21) at both openings. On the narrow opening of 30 kW the minimal cut sets are the three single components, {UA}, {UB}, {feeder}: the tree is one OR gate over three basic events, and the one path set, {UA, UB, feeder}, is the series diagram. On the wide opening of 80 kW they are {feeder}, {UA, UB}: the tree is an OR gate over the feeder and an AND gate over the two units, and the path sets {UA, feeder}, {UB, feeder} are the two branches of the diagram that puts the feeder in series with the parallel pair. These are the artifacts the first process asks the analyst to draw, here derived from the physics, and evaluating either tree by the algebra of Section 3.7 returns the service reliabilities of Table 4. Along the sweep the tree keeps the first shape up to 50 kW and the second from 60 kW on, so the step of Figure 6 is the capacity at which the derived tree changes shape, and nothing else about the tree changes. The tree therefore carries its validity condition: the tree with the AND gate is the system's fault tree for an opening of at least 60 kW and not below, a statement a drawn tree does not make about itself. For the owner, the derived tree is the artifact a standard asks for, with the condition under which it holds written beside it.

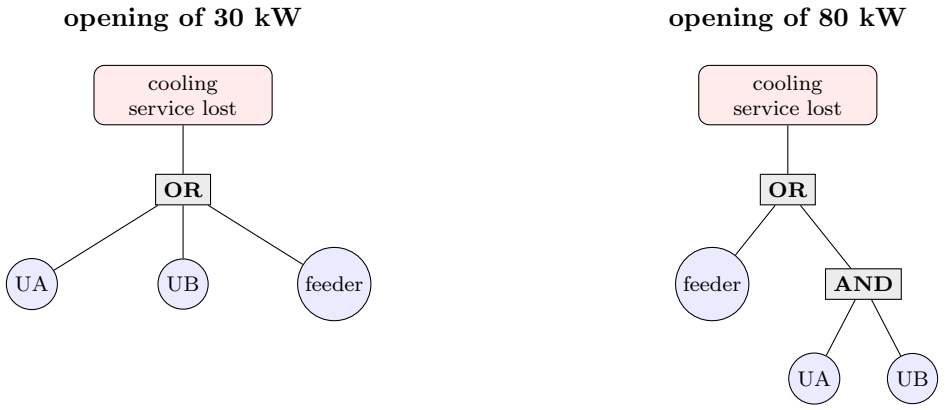


Figure 8: The fault tree read off the enumeration at each opening, by (22): the top event is the OR of one gate per minimal cut set of Table 5, a single-component cut set entering as a basic event and a larger one as an AND gate over its components. On the narrow opening the tree has three basic events under the OR gate; on the wide one the two units sit under an AND gate beside the feeder. The shape changes at an opening of 60 kW and nowhere else along the sweep.

Table 6: The common-cause-failure variant on the wide opening under boundary II. The product of the averaged reliabilities exceeds the service reliability, and the cooling-equipment reliability conditional on the feeder being up is below its unconditional value: the product form (18) no longer holds.

quantity	value [–]
service reliability under the common-cause failure, (23)	0.9725
feeder reliability averaged over the cause, R_F	0.9830
cooling-equipment reliability averaged over the cause, R_C	0.9944
product of the averaged reliabilities, $R_F R_C$	0.9775
cooling-equipment reliability given the feeder is up, $R_{C\|F}$ by (24)	0.9893

4.5 Common-cause failure

Table 6 gives the common-cause-failure variant of Section 3.9 on the wide opening under boundary II. The service reliability by (23) is 0.9725. The feeder's and the pair's reliabilities averaged over the cause are 0.9830 and 0.9944, and their product, 0.9775, exceeds the service reliability by 0.0050. The cooling-equipment reliability conditional on the feeder being up, by (24), is 0.9893, below the unconditional 0.9944: the disturbance state is where both the feeder and the units are more likely to be out, so knowing the feeder is up makes the units more likely to be up as well, and the product of averaged reliabilities, which ignores that, overstates the service reliability. The direction is general: a common-cause failure of the kind modeled here moves the conditional reliability below the unconditional one. For the owner, a product of reliabilities is a claim that the parts fail independently, and a shared cause is where that claim fails first.

5 Discussion

The comparison in this report is between two processes, and the numbers serve only to show where they part. The first process is not wrong: both of its assertions are exact for the structure they assume, and on this system each is the truth at one opening. What the first process lacks is a way to tell which assertion applies, because the quantity that decides it, the heat-flow capacity of

the opening, is not a component and appears on no equipment list. The second process carries that quantity as a relation of the network, (3), and the verdict of each state follows from it. Three further differences follow from the same construction. The boundary of the study is the list of components enumerated, so an assumption such as the availability of the electrical supply is a visible absence from a list rather than a sentence in a report. The independence of the components is a stated property of (11), and relaxing it to the mixture (23) changes the result without changing the physics; whether a common-cause failure exists, and what weight it carries, remain inputs that the analyst supplies, and the report does not discover them. And the enumeration carries its own bound (12) on what it leaves out. That bound is an identity, one minus the listed mass, and this report exercises it only in the trivial case of a complete list, where it is round-off; its value on a system too large to enumerate in full is not demonstrated here.

The second process is not an alternative to the first in what it consumes or in what it claims. It takes the same component list and the same failure data, the unavailabilities of Table 1, from the same failure-mode analysis and the same handbooks a block diagram draws on, and changes neither; where the diagram's structure is right it returns the diagram's number, as Section 4.2 shows at both openings. Nor does it produce a different redundancy count: an $N + 1$ claim is a statement about the topology, and the enumeration checks whether that statement holds once the physics between the components is included, naming the quantity that decides it. Nor does it give up the artifacts of the first: the minimal cut sets and path sets of Section 4.4 are read off the same enumeration, so the one-page fault tree, its closed-form evaluation and the block diagram a standard asks for are outputs of the second process, each with the condition under which it holds. The cost of the second process is that its judgments move rather than vanish. The thermal network, the re-share rule (5) and the support map of Section 3.6 are modeling choices, and a wrong one is wrong as silently as a wrong diagram; what changes is that the choices are relations on a network, which can be checked against a measurement or a more detailed model, rather than a drawing that encodes a conclusion. The limits of this study are those choices. The physics is steady and linear, with no thermal mass, so the time a hall takes to overheat after a loss does not enter, and the limit checked is a heat-flow capacity rather than a hall temperature. The unavailabilities are hour-level availabilities, and the service reliability is the probability that the service is delivered at a random hour, not the probability of surviving a mission of stated length with repair, which is a different calculation. The feeder's effect is declared, not solved: a feeder with a transfer switch, two feeders, or a battery-backed supply would each change the support map, and an electrical network coupled to the units' capacities would compute it. The system has two to three components by design, so that every state could be listed and every reliability checked against its closed form; the enumeration, the scoring and the mixture are the same on a system with hundreds of components, where the structure cannot be drawn by hand and the bound (12) is what bounds a truncated list.

The capacity C deserves a physical reading, because it is the quantity the first process cannot carry and the one this report sweeps. For an opening that carries air it is the product of two things, the airflow the opening can pass and the temperature rise the halls can tolerate while it does,

$$C = \dot{m}_{\max}\, c_p\, \Delta T_{\text{allow}}, \tag{28}$$

with $\dot{m}_{\max}$ the largest mass flow of air the opening can pass, c_p the specific heat of air and ΔT_{allow} the hall temperature rise the IT equipment tolerates; through the opening's conductance the same quantity is $G\,\Delta T_{\text{allow}}$, which is how (4) sets each opening's conductance. Whatever reduces the airflow, the allowable rise, or the mixing that turns the one into the other reduces C, and (28) is the list of ways it changes. Some are the slow wear-out a maintenance program already tracks for the cooling units and seldom for an opening: dust on grilles, louvers and mesh, loaded filters and fouled impellers on any transfer fan, belt and bearing wear and motor derating, each of which

shaves the airflow over months, and blanking panels, floor tiles and containment that drift until the same airflow transfers less heat. Others are discrete, and matter more in practice: a fire or smoke damper that trips, a door that used to be propped open and now closes, pallets, cable trays or a new rack row placed in front of the opening, a partition added for fire separation or acoustics, or new IT equipment with a lower inlet limit, which lowers the allowable rise with no physical change to the opening at all. Two further effects behave like deterioration without being it. With no transfer fan, the flow through an opening is driven by the pressure difference the cooling-unit fans create, so the capacity available in an outage state is not the capacity measured with both units running; this model holds C fixed across states, and a real opening does not. And C is seldom measured: it is a commissioning calculation or a design assumption, and the change a plant sees is often the gap between the as-designed and the as-built value, discovered the first time it matters. None of these mechanisms has anywhere to appear in a block diagram. In the second process each is either a drift in a parameter to re-check, which argues for measuring C periodically by tripping a unit on purpose and reading the hall difference, or a discrete state to enumerate, which argues for treating the opening as a component with its own failure modes, a damper, a door and an obstruction among them.

Table 7 collects the answers of the report.

Table 7: The answers of this report, one row per question a user of the system would ask, in the order of the Conclusion. Every entry is read from the committed results file.

question	answer	why, with the deciding number	what would change it
Are the two units a redundant pair?	conditional, on the opening	with one unit out the service is delivered only when $C \geq 60$ kW, one hall's load (Figure 6)	the heat-flow capacity of the opening, or a transfer duct
What is the service reliability of the pair?	conditional, on the opening	0.9025 below 60 kW and 0.9975 from 60 kW, by (10)	the same
Does either block-diagram assertion give the right number?	yes, each at its own opening	the enumerated reliability equals (16) on the narrow opening and (17) on the wide one, to below 10^{-12} by (27)	nothing identified
Does moving the feeder inside the boundary change the answer?	yes	0.9975 becomes 0.9875 on the wide opening, the feeder's factor	nothing identified
Does the product of feeder and equipment reliabilities hold?	conditional, on independence	it holds to below 10^{-12} with independent components and overstates the service reliability by 0.0050 under the common-cause failure (Table 6)	a common-cause failure of supply and equipment
Can the result be audited as a fault tree or a block diagram?	yes, derived	the minimal cut sets are {feeder}, {UA, UB} on the wide opening and the three single components on the narrow one (Table 5); the tree changes shape at 60 kW	nothing identified
Is the feeder's effect on the units solved from an electrical network?	no	the dependency is declared as a support map (Section 3.6); no electrical network is solved	an electrical network coupled to the units' capacities

6 Conclusion

This report asked whether two cooling units that could each carry both halls form a redundant pair, and what the reliability of the cooling service is; its aim was to compare two processes, not two numbers. The first process asserts the success structure and applies exact algebra to it. The second enumerates every outage state, scores each with the engine's steady heat-flow solution on

the conservation network of the two halls, and sums the states that deliver, so the structure is computed. The engine's part was the network itself, one flow solve per state, and the probability layer that lists states in order with a bound on what it leaves out. What the comparison found: the structure is decided by a quantity no equipment list carries, the opening's capacity, and each block-diagram assertion is right at its own opening; moving the feeder inside the boundary returns the product rule as a result with its conditions stated; a common-cause failure makes that product overstate; and the minimal cut sets read off the enumeration give the fault tree back, with the condition under which it holds. Table 7 holds the numbers. The feeder's effect is declared rather than solved, which is the next step: an electrical network coupled to the units' capacities, on which a transfer switch or a second feeder is a component rather than a map.

A Reproduction

```
cd docs/terra_graph_engine/case_studies/foundation_hvac/two_hall_service_reliability
PYTHONPATH=../../../../../src python run.py
PYTHONPATH=../../../../../src python make_report.py --pdf
```

`run.py` builds the thermal network and the component models, enumerates every outage state, scores each with the engine's steady heat-flow solution, reads the minimal cut sets, path sets and fault tree off each enumeration, and writes `results.json`, with probabilities snapped to a grid of 1e-09 and round-off quantities stated as bounds; `make_report.py` reads only `results.json`. The enumeration is exhaustive (mass target one), so no seed enters. The gate test under `tests/case_studies/foundation_hvac/` pins the closed forms, the step of the sweep, the direction of the common-cause effect, the committed results file and the byte-identical report source.

B Data and problems

The system is the study's own construction and uses no external data. The unavailabilities are design numbers chosen for the demonstration, not published failure data for any product. The study leaves no corpus problems behind; its four cases and its sweep are reproduced in full by the recipe above.

C A worked state

This appendix carries out Section 3.5 by hand for one outage state, the loss of unit B, on both openings, and then runs the other three states in a line each, so that the sum of (10) can be checked against the closed forms. Every number is read from the committed results and the inputs of Table 1; Table 8 collects the four states.

C.1 The state and its probability

Under boundary I the component list is the two units, so the state is $S = \{\mathrm{UB}\}$. By (11) its probability is the product of the unavailability of the unit that is out and the availability of the

one that is not, and both units carry the same unavailability q of Table 1, $q_{\mathrm{UA}} = q_{\mathrm{UB}} = q$,

$$\begin{aligned} P(\{\mathrm{UB}\}) &= q_{\mathrm{UB}}\,(1 - q_{\mathrm{UA}}), \\ &= q\,(1 - q), \\ &= 0.05 \times 0.95, \\ &= 0.0475. \end{aligned}$$

No support is in play: the feeder is outside boundary I, so the units out of service are exactly the units in S.

C.2 The duty the survivor takes up

The failed unit dropped its operating-point duty, so the lost duty is $L = c_B^0 = 60$ kW. The one survivor's headroom is its capacity less its operating-point duty,

$$\begin{aligned} H &= c_A^{\max} - c_A^0, \\ &= 130 - 60 \text{ kW}, \\ &= 70 \text{ kW}. \end{aligned}$$

By (5) the share of that headroom the survivor takes up is

$$\begin{aligned} \alpha &= \min(L, H)/H, \\ &= \min(60, 70)/70, \\ &= 0.857, \end{aligned}$$

so by (6) the survivor runs at

$$\begin{aligned} c_A &= c_A^0 + \alpha\,(c_A^{\max} - c_A^0), \\ &= 60 + 0.857 \times 70 \text{ kW}, \\ &= 120 \text{ kW}, \end{aligned}$$

and by (7) the shortfall is

$$\begin{aligned} D &= \max(L - H, 0), \\ &= \max(60 - 70, 0), \\ &= 0 : \end{aligned}$$

everything the failed unit dropped has been picked up.

C.3 The network

With $c_B = 0$, the heat balance (1) at hall B reads

$$60 - 0 - q_{\mathrm{wall}} = 0,$$

so $q_{\mathrm{wall}} = 60$ kW crosses the opening from B to A; hall B's whole load must leave through the opening because nothing else can take it. The balance at hall A closes,

$$60 - 120 + 60 = 0.$$

The closure law (2) then gives the temperature difference the flow needs. On the narrow opening, whose conductance by (4) is 3 750 W/K,

$$\begin{aligned} T_B - T_A &= q_{\text{wall}}/G, \\ &= 60\,000/3\,750, \\ &= 16 \text{ K}, \end{aligned}$$

and on the wide one, 10 000 W/K,

$$\begin{aligned} T_B - T_A &= 60\,000/10\,000, \\ &= 6 \text{ K}, \end{aligned}$$

against the allowable 8 K: the narrow opening fails the rise exactly where the margin below fails the capacity, which is what (4) arranges.

C.4 The margin and the verdict

The margin (13) is the smaller of the two distances to failure, the headroom left after the re-share, $H - L = 70 - 60 = 10$ kW, and the capacity left on the opening, $C - |q_{\text{wall}}|$. On the narrow opening, $C = 30$ kW, it is

$$\begin{aligned} g &= \min\left(H - L,\ C - |q_{\text{wall}}|\right), \\ &= \min(10,\ 30 - 60), \\ &= -30 \text{ kW}, \end{aligned}$$

negative, so the indicator (14) is 0 and the service is lost: the opening cannot carry hall B's load. On the wide opening, $C = 80$ kW, it is

$$\begin{aligned} g &= \min(10,\ 80 - 60), \\ &= 10 \text{ kW}, \end{aligned}$$

non-negative, so the indicator is 1 and the service is delivered. Nothing was asserted: the same state is lost at one opening and delivered at the other because the capacity entered the margin, and that is the whole difference between the two processes.

C.5 The other three states

With nothing out, $S = \{\}$, both factors of (11) are availabilities,

$$\begin{aligned} P(\{\}) &= (1 - q_{\text{UA}})\,(1 - q_{\text{UB}}), \\ &= (1 - q)^2, \\ &= 0.9025, \end{aligned}$$

and $L = 0$: each unit runs at its own 60 kW, no heat crosses the opening, and the survivors keep $H - L = 140$ kW of headroom, so $g = C$ on either opening and the state delivers. With unit A out,

$$\begin{aligned} P(\{\text{UA}\}) &= q_{\text{UA}}\,(1 - q_{\text{UB}}), \\ &= q\,(1 - q), \\ &= 0.0475, \end{aligned}$$

Table 8: The four outage states of boundary I carried through Section 3.5 by hand: the probability (11), the lost duty and headroom, the share (5), the unit duties (6), the heat on the opening from (1), the shortfall (7), the margin (13) and the verdict (14) on each opening.

state	$P(S)$ [–]	L [kW]	H [kW]	α [–]	c_A, c_B [kW]	q_{wall} [kW]	D [kW]	g at 30 / 80 kW [kW]	verdict at 30 / 80 kW [–]
none	0.9025	0	140	–	60, 60	0	0	30 / 80	delivered / delivered
unit A	0.0475	60	70	0.857	0, 120	60	0	-30 / 10	lost / delivered
unit B	0.0475	60	70	0.857	120, 0	60	0	-30 / 10	lost / delivered
both units	0.0025	120	0	–	0, 0	60	120	-120 / -120	lost / lost

the state is the mirror image of the one above, with 60 kW crossing from A to B; it is lost at 30 kW and delivered at 80 kW. With both units out,

$$\begin{aligned} P(\{\text{UA}, \text{UB}\}) &= q_{\text{UA}}\, q_{\text{UB}}, \\ &= q^2, \\ &= 0.0025, \end{aligned}$$

$L = 120$ kW and no survivor, $H = 0$, so by (7) $D = 120$ kW, $H - L = -120$ kW and $g = -120$ kW on either opening, and the state is lost on the shortfall before the capacity is consulted. Table 3 still lists 60 kW on the opening for this state: the solver books the unserved duty at the reference hall A and therefore still routes hall B's load toward it; the number has no bearing on the verdict.

C.6 The sum over the four states

By (10) the service reliability is the sum of the probabilities of the delivering states, on the narrow opening

$$R_{\text{service}} = 0.9025 \times 1 + 0.0475 \times 0 + 0.0475 \times 0 + 0.0025 \times 0, \tag{29}$$

which comes to 0.9025, the $(1 - q)^2 = 0.9025$ of (16), the series assertion, and on the wide opening

$$R_{\text{service}} = 0.9025 \times 1 + 0.0475 \times 1 + 0.0475 \times 1 + 0.0025 \times 0, \tag{30}$$

which comes to 0.9975, the $1 - q^2 = 0.9975$ of (17), the parallel assertion. The four probabilities sum to one, so the un-enumerated mass (12) is zero. The two closed forms were not supplied; they are what the physics summed to.

D Notation

The symbols used in the report are collected in Table 9.

Table 9: Notation used in the report.

symbol	meaning	unit
L_i	IT heat load of hall i	W
c_i, c_i^0, $c_i^{\max}$	duty of the unit at hall i after an outage, at the operating point, and its capacity	W
L, H	duty dropped by the failed units, headroom of the surviving units	W
T_i	temperature of hall i	K
G, C	conductance and heat-flow capacity of the wall opening	W/K, W
q_{wall}	heat flow on the opening	W
α, D	share of headroom taken up by the surviving units, and duty left unserved	–, W
S, $\bar{S}$	an outage state, the set of components out of service; the same state with both units added when the feeder is in it	–
K, P, U, $\mathcal{K}$	a minimal cut set, a minimal path set, the full component list, the collection of minimal cut sets	–
$K' \subsetneq K$, $S \subseteq S'$, $U \setminus P$	a proper subset of K; S contained in S'; the list with the components of P removed	–
$\mathbf{1}[S]$	the indicator of state S: one when the service is delivered, zero when it is lost	–
$\prod$, $\sum$, $\bigvee$, $\bigwedge$	product and sum over the components or states named under them; OR and AND over them	–
q, q_F	unavailability of a cooling unit and of the feeder	–
$P(S)$	probability of the outage state S	–
$g(S)$	margin of state S; the service is delivered when $g(S) \geq 0$	W
R_{service}, R_F, R_C, $R_{C\|F}$	reliability of the service, the feeder, the cooling equipment, and the equipment given the feeder	–
ε	probability mass of the states not enumerated	–
ρ, H	ratio of unavailabilities, and hours per year out of service	–, h
$\dot{m}_{\max}$, c_p	the largest mass flow of air the opening can pass, and the specific heat of air	kg/s, J/(kg K)
ΔT_{allow}	the hall temperature rise the IT equipment tolerates	K
w, π_w	state of the common cause and its weight	–

References

[Rausand and Hoyland 2004] Rausand, M., Høyland, A. (2004). System Reliability Theory: Models, Statistical Methods, and Applications, 2nd ed. Wiley, Hoboken. ISBN 978-0-471-47133-2.

[Billinton and Allan 1992] Billinton, R., Allan, R. N. (1992). Reliability Evaluation of Engineering Systems: Concepts and Techniques, 2nd ed. Plenum Press, New York. doi:10.1007/978-1-4899-0685-4.

[Vesely et al. 1981] Vesely, W. E., Goldberg, F. F., Roberts, N. H., Haasl, D. F. (1981). Fault Tree Handbook. NUREG-0492, U.S. Nuclear Regulatory Commission, Washington. https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr0492/.

[Ball et al. 1995] Ball, M. O., Colbourn, C. J., Provan, J. S. (1995). Network reliability. In: Ball, M. O., Magnanti, T. L., Monma, C. L., Nemhauser, G. L. (eds.), Network Models, Handbooks in Operations Research and Management Science, vol. 7, chapter 11. Elsevier, Amsterdam.

[Mosleh et al. 1988] Mosleh, A., Fleming, K. N., Parry, G. W., Paula, H. M., Worledge, D. H., Rasmuson, D. M. (1988). Procedures for Treating Common Cause Failures in Safety and Reliability Studies, vol. 1. NUREG/CR-4780, U.S. Nuclear Regulatory Commission, Washington.

[USNRC 1983] U.S. Nuclear Regulatory Commission (1983). PRA Procedures Guide: A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants. NUREG/CR-2300, vols. 1–2. American Nuclear Society and IEEE, Washington.

[IEEE 2013] IEEE (2013). IEEE Std 3006.7-2013, Recommended Practice for Determining the Reliability of 7x24 Continuous Power Systems in Industrial and Commercial Facilities. IEEE, New York. doi:10.1109/IEEESTD.2013.6493367.

[IEEE 2014] IEEE (2014). IEEE Std 3006.5-2014, Recommended Practice for the Use of Probability Methods for Conducting a Reliability Analysis of Industrial and Commercial Power Systems. IEEE, New York.

[IEEE 2018] IEEE (2018). IEEE Std 3006.8-2018, Recommended Practice for Analyzing Reliability Data for Equipment Used in Industrial and Commercial Power Systems. IEEE, New York.

[IEEE 2025] IEEE (2025). IEEE Std 3006.1-2025, Recommended Practice for Reliability Planning and Design of Industrial and Commercial Power Systems. IEEE, New York.