

Periodic readiness testing with repeated detection uncertainty and offline service windows

Mahmoud Abdelfattah Abbas Mahmoud

mahmoud.hti@gmail.com

High Performance Industry (HPI)

Research Article

Keywords: standby generator, dormant failure, periodic inspection, repeated detection, offline service, demand readiness

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11269129/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: Competing interest reported. The author is employed at High Performance Industry (HPI). This employment is disclosed as relevant professional context. No company or client operational records, proprietary equipment data or financial data were used. The author has declared no other competing interests.

Periodic readiness testing with repeated detection uncertainty and offline service windows

Abstract

Shortening a standby test interval reduces hidden fault exposure but increases offline time. This analytical and computational study derives exact stationary demand readiness for a hypothetical single unit with exponential dormant faults, repeated independent restoration opportunities and residual test faults. A strict concavity argument proves one continuous optimum for positive hazard and offline duration. Across 81 declared combinations, optimal intervals range from 1.60 to 22.28 days. The nominal optimum is 6.44 days with readiness 0.9719; the worst uncertainty corner gives 6.69 days and 0.8850. A three state extension distinguishes repeatedly detectable faults from persistent test blind spots and adds a perfect full test every selected number of cycles, with a longer offline window. At a seven day cycle, a 20 percent blind fault fraction and full testing every thirteen cycles give readiness 0.9469, compared with 0.9696 for no blind faults under the same full test schedule. Numerical selection shifts the interval to 4.55 days and readiness to 0.9510. Without full tests, a positive blind fault arrival rate ultimately drives readiness to zero. Independent matrices and 4.68 million additional simulated cycles check the extension. The result quantifies the consequence of the detection assumption; it does not prescribe a regulatory schedule or establish field generator reliability.

Keywords: standby generator; dormant failure; periodic inspection; repeated detection; offline service; demand readiness

1 Introduction

Standby equipment can be apparently available while carrying a fault that is exposed only when it is demanded or inspected. Periodic testing addresses this uncertainty but can itself occupy the equipment and leave residual faults. Jin et al. examined proof test interval and coverage in safety instrumented functions (Jin et al. 2016). Innal et al. showed that partial test formulations require careful treatment of system states and repair times (Innal et al. 2016). These works motivate explicit assumptions about what a test can detect and when restoration occurs. Their safety integrity measures and hardware architectures are not adopted as field parameters here.

Inspection scheduling and standby availability are established research topics. Mendes and Ribeiro optimized inspection intervals for a two unit cold standby system with repair (Mendes and Ribeiro 2017). More recent work by Levitin et al. considered inspection and standby policies under common and operational shocks (Levitin et al. 2025). The present study addresses a narrower operational question: how does a single standby unit's long run readiness change when detection is repeatedly uncertain and each inspection occupies a fixed offline window? It contains no redundant standby component or battery reserve and excludes running mission reliability after a successful demand.

The distinction between a repeated detection chance and a permanently undetectable class is central. A hidden fault that has a fresh independent chance of detection at every test can eventually be restored, even if the probability at one test is small. A fault mode that a particular test cannot detect may persist indefinitely until a different test occurs. These assumptions should not share the same coverage symbol without explanation. This paper calls its parameter a detection and restoration probability, rather than assigning it a standard proof test coverage interpretation.

Zhao and Yun modeled periodic imperfect inspection for a finite-horizon one-shot system, including failures discovered at later inspections (Zhao and Yun 2019). Repeated discovery is therefore an established modeling mechanism. The present contribution combines an exact stationary single-unit state update, a strict concavity argument establishing a unique optimum, scenario comparisons with fixed calendar intervals, and independent numerical and stochastic checks. No novelty claim is made for Markov inspection modeling itself. Brissaud and Oliveira explicitly separate proof-test-detectable and proof-test-undetectable dangerous failures (Brissaud and Oliveira 2015). Their case study neglects testing duration and models exponential restoration. The present

extension retains fixed offline windows and tests the consequence of a persistent blind class with a periodic full restoration window. Neither fault classification nor partial and full testing is claimed as a new general modeling principle.

2 Materials and methods

2.1 Cycle and state definitions

One cycle has length τ days. Its first x equal to τ minus δ days are available for an immediate service demand if the unit is healthy. Its final δ days form an offline test and service window during which the unit cannot serve a demand. A hidden fault arrives during the available part with exponential hazard λ per day. Once faulty, the unit remains unavailable until a later window succeeds in restoring it. No repair outside the window is modeled.

At the end of every window, a preexisting fault is detected and restored with probability c . The same unresolved fault receives a new independent restoration opportunity at the next window. After that action, a fault is left with probability q , independently of the previous state. A healthy unit entering the window remains healthy except for this final fault event. Hidden fault progression during the offline window is absorbed into this discrete end state convention. Restoration duration is not drawn separately: successful diagnosis and restoration must fit inside δ by assumption.

This convention defines two embedded states, healthy and faulty, immediately after the offline window. It excludes fault aging, diagnosis learning, repair queues, persistent test blind spots, fuel availability, start failure conditional on a healthy state, protection logic and failures during a subsequent running mission. A demand is assumed to sample calendar time uniformly and independently; the resulting measure is time averaged readiness, not the probability of surviving an entire outage.

2.2 Exact stationary readiness

Let b equal one minus q , z equal the exponential of minus λx , and h be the probability of a healthy state after a window. The update consists of surviving the available interval, restoring a faulty state with probability c , and avoiding the final test induced fault:

$$h_{\text{next}} = b [c + (1 - c) z h]; \quad x = \tau - \delta; \quad z = \exp(-\lambda x) \quad (1)$$

For $\lambda > 0$, $0 < c \leq 1$, $0 \leq q < 1$ and $x > 0$, this update has a unique stationary healthy probability. Averaging the survival probability over the available interval and assigning zero readiness to the offline window yields:

$$h^* = \frac{b c}{1 - b(1 - c) z}; \quad R = \frac{b c (1 - z)}{\lambda \tau [1 - b(1 - c) z]} \quad (2)$$

The time denominator is the full τ , not only x . Excluding δ from that denominator would overstate demand readiness. When c is one and q is zero, the expression reduces to the integral of an exponential healthy survival probability divided by the full cycle duration. A zero-duration test with increasingly frequent perfect restoration approaches readiness one. Positive δ forces readiness to zero as τ approaches δ from above. The uniqueness proof requires $\lambda > 0$ and $\delta > 0$. At λ equal to zero, the continuous-limit readiness is $[bc/(1 - b(1 - c))]$ times $(1 - \delta/\tau)$, and no finite interior optimum is asserted.

2.3 Uniqueness of the optimal interval

Define a equal to b times one minus c and f of x as one minus z divided by one minus az . Readiness is a positive constant times f of x divided by x plus δ . The first two derivatives are:

$$f'(x) = \frac{\lambda z (1 - a)}{(1 - az)^2}; \quad f''(x) = \frac{-\lambda^2 (1 - a) z (1 + az)}{(1 - az)^3} \quad (3)$$

Because f is increasing and strictly concave for the stated parameter domain, the function g of x equal to x plus δ times f prime minus f has strictly negative derivative. Its value at zero is positive for δ greater than zero, and its limit at infinity is minus one. Therefore exactly one positive root exists, and that root defines the global maximum over τ greater than δ :

$$(x + \delta) f'(x) - f(x) = 0; \quad \tau^* = x^* + \delta \quad (4)$$

The proof applies to this model, not to maintenance schedules with cost discontinuities, fixed permissible dates or mandatory maximum intervals. If a site's authorized test intervals form a discrete or constrained set, the readiness values should be compared only within that set. Optimization here is a sensitivity exercise and does not establish permission to modify a required testing schedule.

2.4 Parameters and uncertainty assessment

All values in Table 1 are scenario assumptions. The reciprocal hazard values are chosen dormant waiting time scales, not mean running times between failures inferred from a generator fleet. They are independent of the aggregate generator reliability values used in separate battery-supported mission models.

Table 1 Declared parameter values

Parameter	Values	Interpretation
lambda	1/180, 1/365, 1/730 per day	Hidden fault arrival hazard during available time
c	0.50, 0.80, 1.00	Fresh detection and restoration chance at each window
q	0, 0.002, 0.010	Residual test induced fault probability per window
delta	0.5, 2, 8 hours	Complete offline test and service window

The full factorial design contains 81 combinations. Each is compared at its continuously optimized interval and at 7, 30 and 90 days. The nominal parameters are lambda equal to 1 over 365 per day, c equal to 0.80, q equal to 0.002 and δ equal to two hours. The uncertainty box extends from 1 over 730 to 1 over 180 per day, 0.50 to 1.00 for c , zero to 0.010 for q and 0.5 to eight hours for δ . Parameters need not have these ranges in an actual installation.

At fixed τ greater than the largest δ , readiness decreases with lambda, q and δ and increases with c . Thus the lowest readiness in the specified box is at the highest lambda, q and δ and lowest c . Maximizing readiness at this corner gives a conservative interval for the box. This is a mathematical bound under declared ranges, not a confidence bound from observed maintenance records.

2.5 Computation and independent checks

The calculations use Python, NumPy 2.3.5 and SciPy 1.17.0 (Virtanen et al. 2020). The unique derivative root is solved using a bracketed method. The embedded two state transition matrix is also constructed independently and its stationary distribution solved numerically. A separate script iterates the transition matrix and integrates healthy survival using numerical quadrature. An 18001 point interval grid checks that no grid value exceeds the analytical optimum in any of the 81 scenarios.

Four stochastic checks simulate the event chronology for 400000 retained cycles per case after a 5000 cycle warmup, for 1.6 million retained cycles overall. Healthy lifetime, detection and test induced fault outcomes are drawn separately. Each estimate uses 80 batches of 5000 cycles for its standard error. Approximate 95 percent Monte Carlo intervals use the Student t quantile for 79 degrees of freedom; these quantify simulation error, not uncertainty in lambda, c , q or δ . Online Resource 1 supplies inputs, seeds and outcomes. Additional review

checks recover the perfect-restoration baseline, the zero-hazard limit and the vanishing available-window limit, and independently check the optimum equation.

2.6 Persistent blind faults and periodic full testing

A sensitivity extension has states H for healthy, D for a fault detectable by the ordinary test, and B for a fault that the ordinary test cannot detect. A newly occurring dormant fault belongs to B with probability p and to D otherwise. The total hazard remains λ ; fault classes are mutually exclusive and absorbing during the available interval. An ordinary window restores D with probability c and never restores B. After a healthy or successfully restored state, a test induced fault occurs with probability q and is classified by the same p . An unresolved prior fault keeps its class. With z defined for that window, the available-phase transition from H is $(z, (1 - z)(1 - p), (1 - z)p)$, while D and B persist.

Table 2 specifies the ordinary-window matrix in state order H D B. All transitions include the residual fault convention. Every m th window is instead a perfect full test that restores both fault classes before the same residual fault mechanism; all its matrix rows are then $(b, q(1 - p), qp)$. Perfect full restoration is an optimistic scenario assumption, not an inspection capability inferred from equipment records.

Table 2 Ordinary-window transitions

Entering state	Leaving H	Leaving D	Leaving B
H	b	$q(1 - p)$	qp
D	cb	$1 - c + cq(1 - p)$	cqp
B	0	0	1

The cycle starts after a full window with probabilities $(b, q(1 - p), qp)$. Propagation through m cycles gives the healthy start probability h_k for each cycle. Its averaged readiness is:

$$R_m = \frac{\sum_{k=1}^m h_k [1 - \exp(-\lambda(\tau - \delta_k))]}{m \lambda \tau} \quad (5)$$

The first m minus one offline windows last two hours; the final full window lasts eight hours. The cycle length τ is common and greater than eight hours. Nominal λ , c and q are retained. The factorial sensitivity uses p equal to zero, 0.05 or 0.20 and m equal to 4, 13 or 52. A bounded numerical interval search from just above eight hours to 365 days is checked against a separate interval grid; no uniqueness proof is asserted for this extension. A second implementation solves the stationary distribution of the product of all phase matrices. Independent full-test cycles permit regenerative simulation standard errors. Three checks retain 4.68 million test cycles in total.

With no full test, positive λ and p make B accessible and absorbing. Consequently the limiting stationary readiness is zero, even if ordinary tests are frequent. At p equal to zero and with ordinary windows only, the original two state stationary update is recovered. These limiting cases explain why replacing a permanent coverage fraction by repeated chance c can be structurally optimistic.

2.7 AI assistance in study preparation

OpenAI ChatGPT, accessed on 4 and 5 October 2026, assisted substantially with topic development, model formulation, mathematical reasoning, software, numerical verification, literature discovery, drafting, figures from computed outputs and Arabic translation. All reported simulations were executed by the supplied code. AI is not an author and does not replace human accountability. Further assistance on 5 October 2026 covered journal-specific formatting, reference presentation and submission preparation.

3 Results

3.1 Interval and readiness comparisons

Optimal intervals across the 81 declared scenarios range from 1.601 to 22.283 days. The nominal optimum is 6.4419 days, with readiness 0.971885. The conservative box optimum is 6.6871 days, with worst case readiness 0.885005. Figure 1 plots both curves and their optima. The two optimized intervals are close in this particular example; a highly precise continuous interval gives little practical advantage over a seven day comparison.




Fig. 1 Long run demand readiness versus test interval in the nominal and worst corner scenarios. Markers identify the unique continuous optima. A shorter interval can increase offline exposure while a longer interval increases hidden fault exposure

Table 3 separates nominal and conservative readiness. The worst case readiness at the nominal optimum is 0.884940, only 0.000065 below the conservative optimum. At seven days it is 0.884908, whereas thirty and ninety day intervals give 0.777797 and 0.555514. These differences are conditional changes in model inputs and scheduling, not estimates of an observed causal maintenance effect.

Table 3 Selected schedules

Schedule	Interval days	Nominal readiness	Worst corner readiness
Nominal optimum	6.4419	0.971885	0.884940
Conservative optimum	6.6871	0.971867	0.885005
Seven day comparison	7.0000	0.971797	0.884908
Thirty day comparison	30.0000	0.936674	0.777797
Ninety day comparison	90.0000	0.837701	0.555514

Figure 2 varies the offline window for a fixed hazard and residual test fault probability. Longer offline windows shift the optimum toward longer intervals. A larger restoration chance permits a different balance between exposure accumulated during standby and time spent offline. The curves describe sensitivity to assumed inspection performance, not measured test effectiveness.




Fig. 2 Optimal interval as a function of the offline window at lambda equal to 1 over 365 per day and q equal to 0.002. Each curve uses the indicated repeated detection and restoration probability

3.2 Numerical and stochastic verification

The largest difference between the exact readiness and the independently solved stationary matrix is 4.44 times ten to the minus sixteen. The independent quadrature verifier differs by at most 5.55 times ten to the minus sixteen. No dense grid point exceeds its analytical optimum. Enumeration of all 16 uncertainty corners at 181 test intervals agrees with the worst corner rule exactly in the saved arithmetic.

Table 4 reports the chronology checks. All four exact values fall within their estimated 95 percent simulation intervals. The largest standardized discrepancy is 1.56 standard errors. Sampling agreement checks the state and timing implementation under the same assumptions and does not establish their suitability for a real generator.

Table 4 Independent chronology checks

Case	Exact readiness	Simulated readiness	95 percent half width	Standardized difference
Nominal at seven days	0.971797	0.971719	0.000429	minus 0.36
Nominal at thirty days	0.936674	0.937304	0.000806	1.56
Worst corner at optimum	0.885005	0.884665	0.001164	minus 0.58
Worst corner at ninety days	0.555514	0.555211	0.002243	minus 0.27

3.3 Consequences of a persistent blind class

For p equal to 0.20 and a full test every thirteen cycles, weekly readiness is 0.946858, compared with 0.969602 for p equal to zero under the same full-test timing. The selected interval becomes 4.5483 days with readiness 0.950975. Comparison with the original 6.4419 day optimum also changes full-test timing, so Table 5 provides the more controlled comparison at a common seven day cycle.

Table 5 Twenty percent blind-fault sensitivity

Full test every m cycles	Readiness at 7 days	No blind faults at 7 days	Selected interval days	Readiness at selected interval
4	0.959413	0.964659	7.6557	0.959564
13	0.946858	0.969602	4.5483	0.950975
52	0.876586	0.971248	2.4262	0.917432

Longer spacing of full tests permits more persistent blind-fault exposure. Figure 3 shows the effect together with the offline penalty at short ordinary cycles. The zero-blind comparison also has full tests, so it does not silently remove their longer offline windows.




Fig. 3 Readiness with a persistent blind-fault class and periodic perfect full tests. Ordinary and full offline windows are two and eight hours respectively. The dashed comparison has no blind class and full testing every thirteen cycles

The 36 selected interval comparisons agree with the independent stationary product-matrix calculation within 7.33 times ten to the minus fifteen. The three simulated means are 0.946916, 0.950470 and 0.877484, against exact values 0.946858, 0.950975 and 0.876586. Their standardized differences are 0.14, minus 1.45 and 0.77; all exact values lie within the corresponding estimated 95 percent Monte Carlo intervals. Numerical checks of the ordinary-only zero-blind limit and blind-state absorption also pass.

4 Discussion

The detection assumption materially changes the meaning of the result. An unresolved fault receives a new independent chance c of restoration at every window. This can represent uncertain inspection execution within a simplified operational audit, but does not represent a permanent collection of fault modes that the test cannot reveal. Partial test modeling in the literature (Innal et al. 2016; Jin et al. 2016) must therefore not be replaced automatically by the present formula. The extension makes that distinction explicit and quantifies its effect. The partial/full-test methods of Innal et al. (2016) and the fault classes in (Brissaud and Oliveira 2015) are closer conceptual comparators than a single repeated chance parameter. This is a conditional mechanism comparison, not a claim to outperform either publication's system or reproduce its formulas.

The single unit readiness measure also differs from the two unit repair model of Mendes and Ribeiro (2017) and the common shock standby policies of Levitin et al. (2025). This study obtains a transparent continuous optimum by assuming a fixed offline window and exponential dormant faults. It excludes costs, redundancy, continuous diagnosis, subsequent run failures and shared supply failures. Its conservative optimum describes

the stated rectangular parameter set only. A narrow difference between the nominal, conservative and seven day values should not be exaggerated into an operational benefit.

All successful restoration is assumed to complete within Δ . Actual diagnosis and parts delays may produce long or dependent repair durations, making that assumption optimistic. Fault hazards can depend on age, environment, load testing history and maintenance. The test induced fault parameter may also depend on the test frequency or service quality; fixing it per window is a modeling choice. Demands that arrive more often during selected periods invalidate the uniform independent time average.

The appropriate next validation step is to log hidden faults by discovery time, test duration, unresolved diagnoses, post test failures and requests for service. Those records would support a distinction between repeatedly missed detectable faults and intrinsically undetectable modes. Any practical interval decision would also need the equipment maintenance instructions and the required test regime. The calculated intervals are not a substitute for those requirements.

5 Conclusions

A fixed offline service window and repeated uncertain fault restoration produce an exact periodic readiness function with one continuous optimum. The hypothetical 81 case study shows an optimal interval range of 1.60 to 22.28 days. In the selected nominal and conservative examples, the optima are close to seven days, while thirty and ninety day comparisons lose readiness under the assumed hidden fault hazard. Independent matrices, quadrature and simulated chronologies support the implementation. The three state sensitivity shows that frequent ordinary tests cannot compensate for a persistent blind class without another restoration mechanism. With a declared 20 percent blind fraction and full testing every thirteen cycles, a seven day schedule gives 0.9469 readiness and numerical selection shifts to 4.55 days. Application requires evidence for the detection mechanism and timing assumptions and must respect the site's required maintenance and test regime.

6 Nomenclature

Symbol	Meaning	Unit
$\tau \Delta x$	Cycle offline window and available interval	days
λ	Hidden fault hazard during available interval	per day
c	Repeated detection and restoration probability	dimensionless
q and b	Residual test fault and its complement	dimensionless
z	Probability of no new hidden fault during x	dimensionless
h and h^*	Post window healthy probability and stationary value	dimensionless
R	Long run time averaged demand readiness	dimensionless
$p_m h_k$	Blind-fault fraction full-test spacing and healthy cycle-start probability	dimensionless
$a f g$	Auxiliary quantities in optimality proof	dimensionless

References

Brissaud F, Oliveira LF (2015) Average probability of a dangerous failure on demand: Different modelling methods, similar results. Author-deposited preprint, arXiv:1501.06487. <https://arxiv.org/abs/1501.06487>

- Innal F, Lundteigen MA, Liu Y, Barros A (2016) PFDavg generalized formulas for SIS subject to partial and full periodic tests based on multi-phase Markov models. *Reliability Engineering & System Safety* 150:160–170. <https://doi.org/10.1016/j.ress.2016.01.022>
- Jin J, Pang L, Hu B, Wang X (2016) Impact of proof test interval and coverage on probability of failure of safety instrumented function. *Annals of Nuclear Energy* 87:537–540. <https://doi.org/10.1016/j.anucene.2015.09.028>
- Levitin G, Xing L, Dai Y (2025) Standby and inspection policy optimization in systems exposed to common and operational shock processes. *Reliability Engineering & System Safety* 253:110509. <https://doi.org/10.1016/j.ress.2024.110509>
- Mendes AA, Ribeiro JLD (2017) The establishment of the time interval between inspections for a cold standby system with component repair. *Production* 27:e20170049. <https://doi.org/10.1590/0103-6513.004917>
- Virtanen P, Gommers R, Oliphant TE et al (2020) SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* 17:261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Zhao QQ, Yun WY (2019) Storage availability of one-shot system under periodic inspection considering inspection error. *Reliability Engineering & System Safety* 186:120–133. <https://doi.org/10.1016/j.ress.2019.02.016>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ESM1.zip](#)