Australian Government

Australian Safeguards and Non-Proliferation Office

Australian Safeguards Support Program

Minimum inspection frequencies under Integrated Safeguards

Task AUL C 01208: **Re-Examination of Basic Safeguards Implementation Parameters**

Russell LESLIE, Annette BERRIMAN and John CARLSON

Australian Safeguards and Non-Proliferation Office (ASNO)
RG Casey Bldg, John McEwen Crescent, Barton ACT 0221 Australia

June 2005

AUL Report 2005-01

Page intentionally left blank

SYNOPSIS

Integrated Safeguards (IS) allow the IAEA to adopt the optimum combination of traditional and strengthened safeguards measures in order to achieve overall safeguards goals. Inspections of nuclear facilities are a key traditional measure that will always be a cornerstone of safeguards effort – but in the context of IS it will be possible to take into account a wide range of factors in order to allocate an appropriate level of safeguards resources to facilities of a given type in any state.

This paper examines the issue of the minimum inspection frequency for each state that can be used by the IAEA in order to achieve its goals under IS.

A number of archetypal cases will be considered:

- a state with one nuclear facility;
- a state with more than one nuclear facility but without front- or back- end facilities; and
- a state with a full nuclear fuel cycle.

The paper considers whether it would be possible to arrive at generally applicable minimum inspection frequencies for the three cases considered.

The paper draws on work that has been done by various authors on statistical and game theoretical studies of randomised inspection regimes to produce a method of determining the number of randomised inspections necessary to achieve an adequate level of assurance that the declared activities have been appropriately covered. The required number of inspections would be derived as a function of the chosen detection probability and acceptable non-detection probability. This requires the establishment of a series of parameters relating to inspection goals including a variety of factors relating to non-detection probability. The paper also examines issues such as the verification activities that would be required during these inspections and issues that could affect the effectiveness of such an inspection regime.

PREAMBLE

The level of assurance that can be derived from a given level of inspection effort has been the subject of intensive academic study since the early seventies. The formulae used in the stratification, sampling and verification of nuclear material under safeguards have been examined and reconsidered frequently since they were first derived. The approach taken by the Agency has gained wide acceptance.

This paper assumes the basic theoretical underpinnings of the existing safeguards system. The paper also makes use of a simple Visual Basic program as an independent crosscheck of all results - deriving the same result by both using exponential distribution formulae and a Monte-Carlo style simulation (two independent methods) is used as a means of boosting confidence in the final results.

IMPORTANT NOTE: This paper deals explicitly only with the detection of diversion from the declared inventories of declared facilities. Detection of undeclared activities is beyond the scope of the current paper.

1.1—INTRODUCTION

The possibility of appropriately trained IAEA inspectors being present at nuclear facilities is a key element of the effectiveness of the safeguards system. The presence of inspectors is viewed by many as a powerful deterrent to the use of declared facilities for clandestine purposes as it is not possible for potential proliferators to completely control what the inspectors observe. There are two elements of "inspector presence"-derived assurance to be considered.

Element 1 of the inspection – inspection activities

The first, and most well understood, is the assurance provided by the inspection activities themselves, including:

- examination of the facilities materials accountancy records;
- reconciliation of these records and the reports supplied to the Agency;
- examination of containment and surveillance systems; and
- nuclear materials verification activities (e.g. item counting, serial number identification and measurement).

The deterrent effect of these activities has been the subject of detailed statistical study for the last thirty years with specified detection and non-detection probabilities and detailed evaluation of the quality of the inspection activities performed and formal evaluation of inspection effectiveness.

Element 2 of the inspection – inspector presence

The second element is the deterrence arising from the physical presence of the inspector at the facility.

Under the traditional safeguards system there was an expectation that inspectors would limit their area of interest to strategic points, and there was some resistance to the concept of inspectors exercising their own judgment or curiosity about activities away from these strategic points.

One of the strengthened safeguards measures in broad application is that Agency inspectors receive training on "enhanced observations skills" and are encouraged to exercise a healthy scepticism to the information supplied to them by the facility operators or the state (following the "trust but verify" model), and to be curious about the form and structure of the facility beyond the defined strategic points. The deterrence value of granting access to a facility of a person that has received training on the types of observables that can indicate clandestine activity cannot be quantified but does appear to be highly valued by a broad cross-section of the international community.

There is also a deterrent effect arising from the potential of an inspector at a facility requesting complementary access to areas of the facility that would not normally be seen by an inspector during a standard inspection.

It is important to note that it is possible to devise strategies in which the various parts of element 1 of the inspection could be accomplished without inspector presence - but element 2 of the inspection can be accomplished only by the physical presence of fully trained IAEA inspectors. “Boots on the ground” cannot be replaced by any form of technology or by any other inspection modality.

The discussion in this paper will deal quantitatively with issues relating to element 1 of the inspection but the need to maintain element 2 of the inspection will inform the discussion.

1.2—OPTIMUM TIMING FOR DIVERSION

Key Assumption 1 (KA1): That the diverter is prepared to take all necessary steps to prevent detection (in effect to maximise the time to detection). One key element of any strategy to maximise the time to detection is to divert material immediately after the completion of an inspection.

Key Assumption 1 (KA1) holds throughout the discussion that follows.

1.3—MINIMUM INSPECTION FREQUENCIES - PRELIMINARY CALCULATIONS

In order to examine the concept of minimum inspection frequencies this paper will start with a series of simplified mathematical calculations related to this subject. These calculations make use of the same mathematical tools that are needed to examine realistic inspection issues.

We will make a set of simplifying assumptions for our example calculations - please note that the calculations without these simplifying assumptions are still valid but in some cases adjustments are necessary.

Simplifying Assumption 1 (SA1): The PIV is randomly scheduled as for all other inspections.

Simplifying Assumption 2 (SA2): The detection probability (p) for every inspection is 1 and the non-detection probability (β) for every inspection is 0. This means it is absolutely certain that an inspection will detect any diversion[1].

Simplifying Assumption 3 (SA3): The timeliness goal is 30 days.

Simplifying Assumption 4 (SA4): There is no diversion scenario for the facility in question that can be concealed in less than 24 hours.

Simplifying Assumption 5 (SA5): There are 365 inspection opportunities per year.

With exactly 12 inspections per year chosen at random intervals, the probability of an inspection on any given day is designated p where p=12/365=3.29%. The probability of achieving a timeliness goal of 30 days (or less) is 63%[2] and the mean interval between inspections is 30.4 days.

---

[1] The effect of varying SA2 is discussed in a separate paper AUL Report 2005-02.

[2] The inspection intervals approximate an exponential distribution and the 63% figure and subsequent figures in Table 1 are derived using the cumulative exponential distribution function $F(x) = 1- e^{(-x/\mu)}$ where μ = the assumed mean time to detection. The effect of using the exponential distribution function instead of the binomial distribution is discussed in AUL Report 2005-02.

Table 1 - No fixed PIV

Target number of inspections/year	Probability of an inspection on any day	Probability of achieving TG <= 30 days	Probability of achieving TG <= 15 days	Mean time to detection[3] (days)
12[4]	1 or 0 as per schedule	100%	0	30.4
12	3.29%	63%	39%	30.4
8	2.19%	48%	28%	45.6
6	1.64%	39%	22%	60.8
4	1.10%	28%	15%	91.3
3	0.82%	22%	12%	121.7
2	0.55%	15%	8%	182.4
1	0.27%	8%	4%	365.0

We will now look at the effect on the above figures of varying SA1 (i.e. conducting an annual PIV at the first inspection opportunity of every year).

Table 2 - Fixed PIV (the case when SA1 is varied)[5]

Target number of inspections/year	Probability of an inspection on any day	Probability of achieving TG <= 30 days	Probability of achieving TG <= 15 days	Mean time to detection (days)
12[6]	1 or 0 as per schedule	100%	0	30.4
12	3.28%	63%	39%	30.5
8	2.19%	48%	27%	45.7
6	1.64%	38%	21%	61.0
4	1.09%	27%	15%	91.4
3	0.82%	20%	11%	121.9
2	0.55%	12%	7%	182.7
1	1 or 0 as per schedule	0.0%	0%	365.0

As can be seen from Tables 1 and 2, the effect on the results of the calculation of SA1 is relatively minor (and only significant for a target inspection frequency of four inspections per year or less).

---

[3] Assuming random inspections for an "infinite" number of years – the longer the time period the more closely the distribution conforms to the exponential distribution.

[4] Limiting case - 12 inspections per year at fixed inspection intervals of 30 days.

[5] There is a small, but potentially significant, deviation from the exponential distribution when PIVs are fixed in time rather than at random times, this is most clearly apparent in cases where the target number of inspections per year is small (4 or less per year).

[6] Limiting case - 12 inspections per year at fixed inspection intervals of 30 days.

1.4—KEY RESULTS

Key Result 1 (KR1): The net effect of the foregoing is to arrive at the trivial conclusion (KR1) that: The mean time to detection is inversely proportional to the inspection frequency (e.g. 12 random inspections (RI)/year gives a mean time to detection of 365/12 ~ 30 days).

IMPORTANT NOTE: KR1 is dependent on SA2 - the effect of lower detection probabilities is explored further in AUL Report 2005-02.

Key Result 2 (KR2): For random inspections, the inspection interval will be shorter than the timeliness goal for a proportion of the inspections, and this proportion is dependent on the target number of inspections for the year.

A corollary to KR2 is that, for random inspections, there is a well-defined probability of achieving a given timeliness goal for a given number of inspections per year.

Key Result 3 (KR3): The minimum inspection frequency is dependent on the acceptable probability of achieving the timeliness goal.

2—DISCUSSION

2.1—CASE 1 (C1)

State with one safeguardable nuclear facility (assume RRCA or power reactor), no front or back-end activities and no other nuclear activities

Base Assumption 1 (BA1): The facility is not an on-line refuelled reactor (OLR).

Base Assumption 2 (BA2): The facility has no unirradiated direct use material (UDU).

Base Assumption 3 (BA3): An annual PIV is required and scheduled based on existing timeliness criteria (with the time period between consecutive PIVs being no more than 14 months).

Base Assumption 4 (BA4): Under integrated safeguards the timeliness criteria for spent fuel has been set at 12 months.

Base Assumption 5 (BA5): Facility has greater than one SQ of nuclear material.

Unless one of these five base assumptions is varied, we arrive at situation where only the annual PIV inspections are required to meet Agency inspection goals for both quantity and timeliness, and the minimum inspection frequency is the period between PIV inspections.

2.1.1—VARYING BASE ASSUMPTIONS

Vary BA1 (i.e. assume that the reactor in question is an OLR) with BA2-BA4 unchanged

The Agency has a draft safeguards approach for OLRs under IS that is heavily dependent on verification of the flow of the fuel through the reactor by unattended monitoring equipment. It is the understanding of the authors that the current IS approach for OLRs requires an annual PIV in addition to flow verification. If the requirement for an annual PIV were to change, then the minimum inspection frequency for OLRs would be set by the maintenance/media change schedule for the flow verification equipment. If remote data transmission is assumed, then site visits could, in principle, be reduced to the frequency of reactive or preventive maintenance visits.

Vary BA2 (i.e. assume that there is UDU at the facility) with BA1 and BA3-BA4 unchanged

There is a range of facility types that are caught by this change of base assumptions - UDU is commonly found as:

- HEU at research reactors;
- either HEU or Pu at critical assemblies; or
- Pu (in the form of MOX) at power reactors.

In the case of research reactors and critical assemblies (RRCAs) the Agency categorises these facilities into groups based on their inventories of UDU and other material types. The groupings can be summarised as follows:

Group (I)	reactors with more than 1 SQ of UDU;
Group (II)	reactors with less than 1 SQ of UDU, and more than 1 SQ of any other material type;
Group (III)	reactors with more than 0.5 SQ of UDU and less than 1 SQ of any material type; and
Group (IV)	reactors with less than 0.5 SQ of each and every material type.

Separate considerations apply to RRCAs with a thermal power greater than 25MWth - but these considerations are not relevant to this discussion. Group (IV) reactors have less than 1 SQ of material and are not relevant to this discussion.

Under IS, Group (II) and Group (III) are currently subject only to annual PIVs (due to the change in spent fuel timeliness requirement from 3 to 12 months). In the case of the IS approach that was adopted for Australia (with a Group (II) reactor) it was decided to supplement the annual PIVs with an average of one short notice random inspection (SNRI) per year.

The authors understand that the current IS approach for Group (I) reactors maintains a timeliness goal of one month. Under the traditional approach to safeguards this would require 11 inspections for timeliness purposes in addition to the annual PIV inspection. Relying on KR2 (that is, there is a well-defined probability of achieving a given timeliness goal for any given number of RIs) it is possible to replace the 11 required timeliness inspections with a smaller number of RIs. This approach will be suitable only if it is possible to accept less than 100% probability of achieving existing timeliness goals.

It is possible to realize a 50% probability of achieving a 30-day timeliness goal with 1 PIV plus 7 RI per year (see table 1 for details[7]). A 20% probability requires 1 PIV plus 2 RI/year. The minimum inspection frequency is therefore dependent on the required detection probability.

For power reactors similar considerations apply, but the required timeliness goal for fresh MOX at these facilities is three months (rather than one month for other types of UDU). Table 3 shows the results for differing inspection frequencies.

Table 3 - Fresh MOX - 90-day timeliness goal

Target number of inspections/year	Probability of an inspection on any day	Probability of achieving TG <= 90 days	Probability of achieving TG <= 45 days	Mean time to detection (days)
4[8]	1 or 0 as per schedule	100%	0	91
4	1.10%	63%	39%	91
3	0.82%	52%	27%	122
2	0.55%	39%	21%	182
1	0.27%	22%	15%	365

---

[7] These calculations are dependent on SA2 (p=1, ß=0) and would need to be recalculated to take into account actual target detection probabilities. The authors have produced a simple application to perform these calculations – see paper AUL Report – 2005-2

[8] Limiting case - 4 inspections per year at fixed inspection intervals of 90 days.

2.2—CASE 2 (C2)

State with more than one nuclear facility (assume mixture of RRCA and/or power reactor), no front or back-end activities and no other nuclear activities

Restarting our discussion with a new set of base assumptions:

Base Assumption 1 (BA1): The inspections across all facilities within the state are independent of each other (*inter alia,* this explicitly ignores the effect of "borrowing" scenarios).

Base Assumption 2 (BA2): All of the facilities in the state are identical.

Base Assumption 3 (BA3): Annual PIVs are required and scheduled based on existing timeliness criteria (with the time period between consecutive PIVs at any facility being no more than 14 months).

Base Assumption 4 (BA4: Under IS the timeliness criteria for spent fuel has been set at 12 months.

Base Assumption 5 (BA5): All facilities have greater than one SQ of nuclear material.

Base Assumption 6 (BA6): All facilities are in an equal state of preparation to be inspected.

The result of these assumptions is that C2 reduces to a series of independent C1 results. The fact that the various facilities are all located within the same state is (in this case) irrelevant to the scheduling of inspections. This example can be viewed as a stylised version of the Agency's traditional safeguards system (except that the traditional system recognises the possibility of "borrowing" and addresses it as part of its safeguards approach for the state).

As per C1 (above), only facilities that contained UDU would require inspections beyond the annual PIV to meet planned timeliness goals under IS.

2.2.1—VARYING BASE ASSUMPTIONS

Varying BA1 (while all other BAs remain unchanged) allows for the consideration of the state as a whole.

As all of the facilities in the state are subject to the possibility of RI there is no need to explicitly deal with borrowing scenarios, as the state would always be subject to the risk that the facility "borrowed from" would be inspected during the period of borrowing.

Another implication of varying BA1 is that the reactors can be viewed as a group[9] (conceptually one way of dealing with this is that it is the combined inventories of all of the facilities in the country that the Agency is seeking to control – in effect the diversion risk is

---

[9] SA2 is also explicitly varied ($p<1$, ß>0) to allow for selection of a sample to represent the population of items available for verification.

pooled). The inventory records for all facilities in a group would have to be available at the time of the inspection (this could be achieved via a mailbox arrangement for inventory records). Stratification would take into account the full inventory of a given material type across all facilities with sample size determined on the basis of the full stratum inventory. Every item in the stratum (regardless of which facility it was located in) would have an equal probability of being selected before the actual facility is selected. Items at the selected facility would be oversampled (when compared to the equivalent C1 case), but the entire stratum would be verified by the one inspection (if there were not enough items available at the selected facility to complete the sampling plan, then it would be necessary to choose more than one facility for inspection in order to verify the whole stratum).

In that case, if a 50% or a 20% probability of achieving a timeliness goal is acceptable for a single facility viewed in isolation, the same probability should be equally acceptable when applied across a group of facilities. The facilities can be pooled for the purposes of planning and carrying out RI – depending on the number of facilities involved in the pool this could result in a substantial saving in Person-Days of Inspection effort (PDI).

One way in which the concept of "pooling" could applied, for inspection planning purposes, is for a target to be set for the number of inspections carried out each year across the pool. This target would be used to determine the probability of an inspection at any facility in the pool on any day. This could then be used to guide inspection allocation effort. One consequence of the pooling concept is that while both the average inspection interval and the mean time to detection could be calculated, the figures would apply to the pool as a whole, not to the individual facilities within the pool (with some facilities within the pool receiving only PIVs and some both PIVs and random inspections).

Varying both BA1 and BA2 (while all other BAs remain unchanged) produces a situation in which more than one facility type exists within the state. While the pool concept might still be applicable it would make little sense to pool facilities with widely different timeliness goals, inventories and/or modes of inspection/operation (e.g. a pool containing both power reactors with fresh MOX and Group (III) RRCAs would make little sense for planning purposes).

Varying BA1, BA2 and BA3 (while all other BAs remain unchanged) carries the pooling concept further with the possibility of pooling being used to allocate which facilities would receive a PIV in any given inspection year. Once again, as noted above, the pooling concept would have utility only if the facilities making up the pool had similar characteristics - a power reactor pool and a separate Group (III) RRCA pool would have greater utility than a single pool containing both facility types.

2.2—CASE 3 (C3)

State with a full nuclear fuel cycle

C3 can basically be seen as a complication on C2 with a wider range of timeliness goals and inspection requirements being present in the state. As for C2, it may be possible to introduce pooling of facilities under IS - with the strong caveat that pooling will make sense only in cases where there are similar characteristics that can be used to generate utility.

One important distinction between C2 and C3 is that C3 provides for an important new element to be added to the pooling concept. Whereas under C2 pooling was applied across facilities with similar safeguards characteristics, under C3 pooling could be extended to all materials with similar safeguards characteristics.

For example:

- UF6 in cylinders and UO_2 powder could each be in separate inspection pools that encompass both the conversion facility and the fuel fabrication plant;
- finished fresh fuel assemblies could be in a single inspection pool that encompasses the product store of the fuel fabrication facility and the fresh fuel store of the power reactor;
- spent reactor fuel could be in a single pool that encompasses both the spent fuel pond of the reactor and the head-end of the reprocessing plant, etc.

Minimum inspection frequencies would, as noted in KR3, be related in direct fashion to the acceptable probability of achieving the stated timeliness goal. For any given acceptable percentage probability it is possible to calculate a minimum inspection frequency using the numerical methods developed by the authors.

3—CONCLUSIONS

The setting of minimum inspection frequencies under IS is an issue with profound resource implications for the Agency. The first step in such a process is to delineate the factors that affect the setting of minimum frequencies, and the second is determining which of these factors are purely technical and which require some form of policy input. Each of the cases examined had distinctive characteristics when viewed in the context of minimum inspection frequencies – in spite of these differing characteristics there were key results that were valid in all cases, namely:

Key Result 1: For random inspections, the inspection interval will be shorter than the timeliness goal for a proportion of the inspections and this proportion is dependent on the target number of inspections for the year.

A corollary to KR1 is that, for random inspections, there is a well-defined probability of achieving a given timeliness goal for a given number of inspections per year.

Key Result 2: The minimum inspection frequency is dependent on the acceptable probability of achieving the timeliness goal.

The IAEA will need to reach credible and defensible policy decisions as to:

- the importance it would place on being able to detect diversions at intervals shorter than the current timeliness goals; and
- the acceptable minimum probability of achieving the applicable timeliness goals.

The actual minimum inspection frequencies are directly dependent on these policy decisions.

BIBLIOGRAPHY:

[1] M J, Canty, G Stein and R Avenhaus, “The Reduced Frequency Unannounced Verification Model: Generalization to Variables Sampling”, Proc. 29th Meeting of the INMM, 1988.

[2] M S Lu, T Teichman, “A Scheme for Randomized Inspections”, 13th ESARDA Symposium on Safeguards and Nuclear Material Management, 1991.

[3] M J Canty “Simulation of Inspection Randomization Models”, 13th ESARDA Symposium on Safeguards and Nuclear Material Management, 1991.

[4] R Avenhaus and M Canty, “Timeliness and Deterrence in IAEA Safeguards”, IAEA-SM-351, IAEA Symposium on International Safeguards, 1997.

[5] R Avenhaus and M Canty, “Compliance Quantified – An introduction to data verification”, Cambridge University Press, 1996.

[6] K Nidaira and T Osabe, “Case Study of Effective Safeguards Approach Using Unannounced/Random Inspection for Uranium Fuels”, Proc. 41st Meeting of the INMM, 2001.

[7] R Avenhaus, “Safeguards Systems Analysis”, Plenum Press, 1986.

[8] R Avenhaus, “Material Accountability - Theory, Verification, Applications”, International Institute for Applied Systems Analysis, 1997.

[9] J.L.Jaech and M.Russell, IAEA Section for Statistical Analysis, “Algorithms to calculate samples sizes for Inspection Sampling Plans”, STR-261, March 1991.

Annex A

Mathematical Concepts

Sampling formulae

The standard form of the equation that is used to determine the number of items (N_0) that needs to be selected from a total population (N) to satisfy required statistically criteria is:

$$N_0 = N^*(1- \beta^{(x/M)}).$$

Where:

N = the number of items in the stratum to be verified - which in this case is defined as the total population of a given type of item available for verification across all facilities in the state;

N_0 = the subset of the population that will need to be verified;

β = the non-detection probability for the inspection;

x = the average amount of material in each item in the stratum;

M = the defined significant quantity of the material (e.g. 8 kg for Pu).

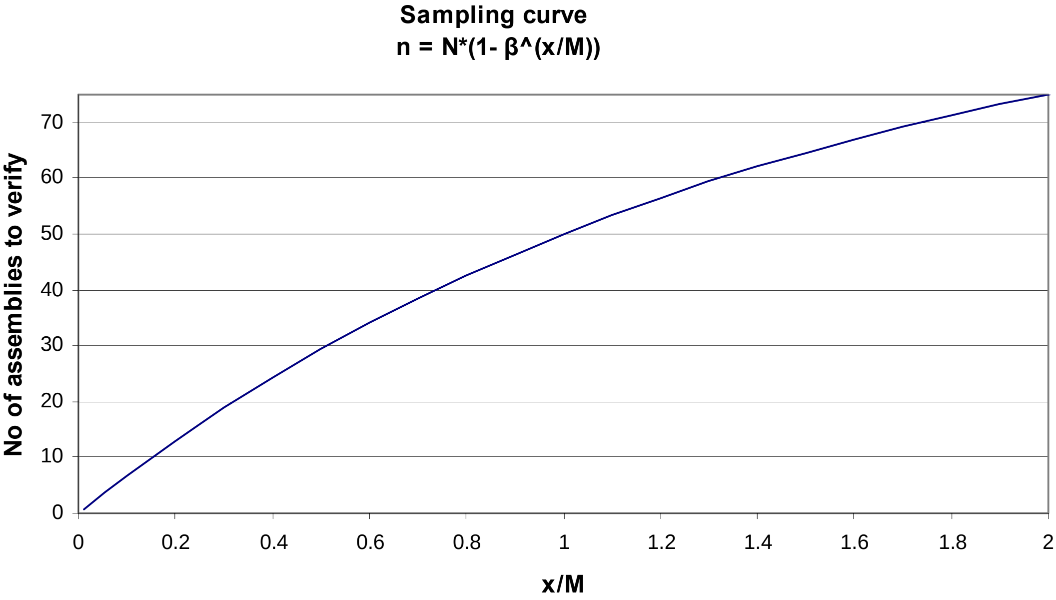

Exponential Distribution.

If we know the population mean (μ) of the times between events that we are interested in, we can define a new exponential function parameter λ that is crucial in doing calculations with the exponential distribution:

$$\lambda = 1/\mu$$
$$= \iota / \psi$$

where ι = the target number of inspections per year and ψ = the number of inspection opportunities per year.

If the mean time to detection is 365/12 = 30.4 days, then

$$\lambda = (12/365) = 0.03288$$

The exponential distribution function is defined as:

$$f(x) = \lambda * e^{-\lambda * x} \quad \text{for } 0 \leq x < \infty, \lambda > 0$$

Once we know the value of λ there are several other factors that can be calculated in straightforward fashion.

$$\begin{aligned} \text{first quartile} &= \ln(^4/_3)/\lambda \\ \text{median} &= \ln(2)/\lambda \\ \text{third quartile} &= \ln(4)/\lambda \\ \text{variance } (\sigma^2) &= 1/\lambda^2 \\ \text{standard deviation } (\sigma) &= 1/\lambda \\ &= \mu \end{aligned}$$

In our case we want to know the probability of achieving our timeliness goal – for that calculation we make use of the cumulative probability distribution equation:

$$F(x) = 1- e^{(-x*\lambda)} \quad \text{for } 0 \leq x < \infty, \lambda > 0$$
$$= 1- e^{(-x/\mu)} \quad \text{for } 0 \leq x < \infty, \mu > 0$$