Dependable Systems

Dependability Attributes

Dr. Peter Tröger

Sources:

J.C. Laprie. Dependability: Basic Concepts and Terminology
Eusgeld, Irene et al.: Dependability Metrics. 4909. Springer Publishing, 2008
Echtle, Klaus: Fehlertoleranzverfahren. Heidelberg, Germany : Springer Verlag, 1990.
Pfister, Gregory F.: High Availability. In: In Search of Clusters. , S. 379-452

Dependability

- **Umbrella term** for **operational** requirements on a system
  - IFIP WG 10.4: "*[..] the trustworthiness of a computing system which allows reliance to be justifiably placed on the service it delivers [..]*"
  - IEC IEV: "*dependability (is) the collective term used to describe the availability performance and its influencing factors : reliability performance, maintainability performance and maintenance support performance*"
  - Laprie: „ *Trustworthiness of a computer system such that reliance can be placed on the service it delivers to the user* “
- Adds a third dimension to system quality
- General question: How to deal with unexpected events ?
- In German: ‚Verlässlichkeit‘ vs. ‚Zuverlässigkeit‘

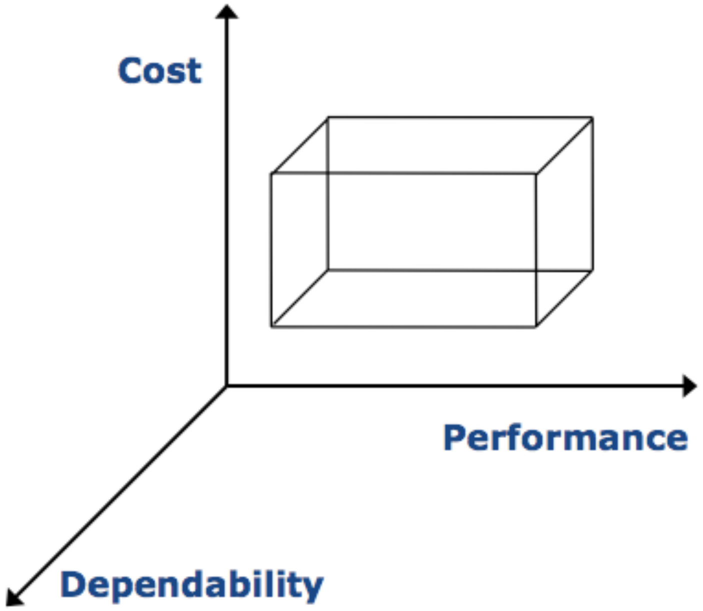

Dependability Tree (Laprie)

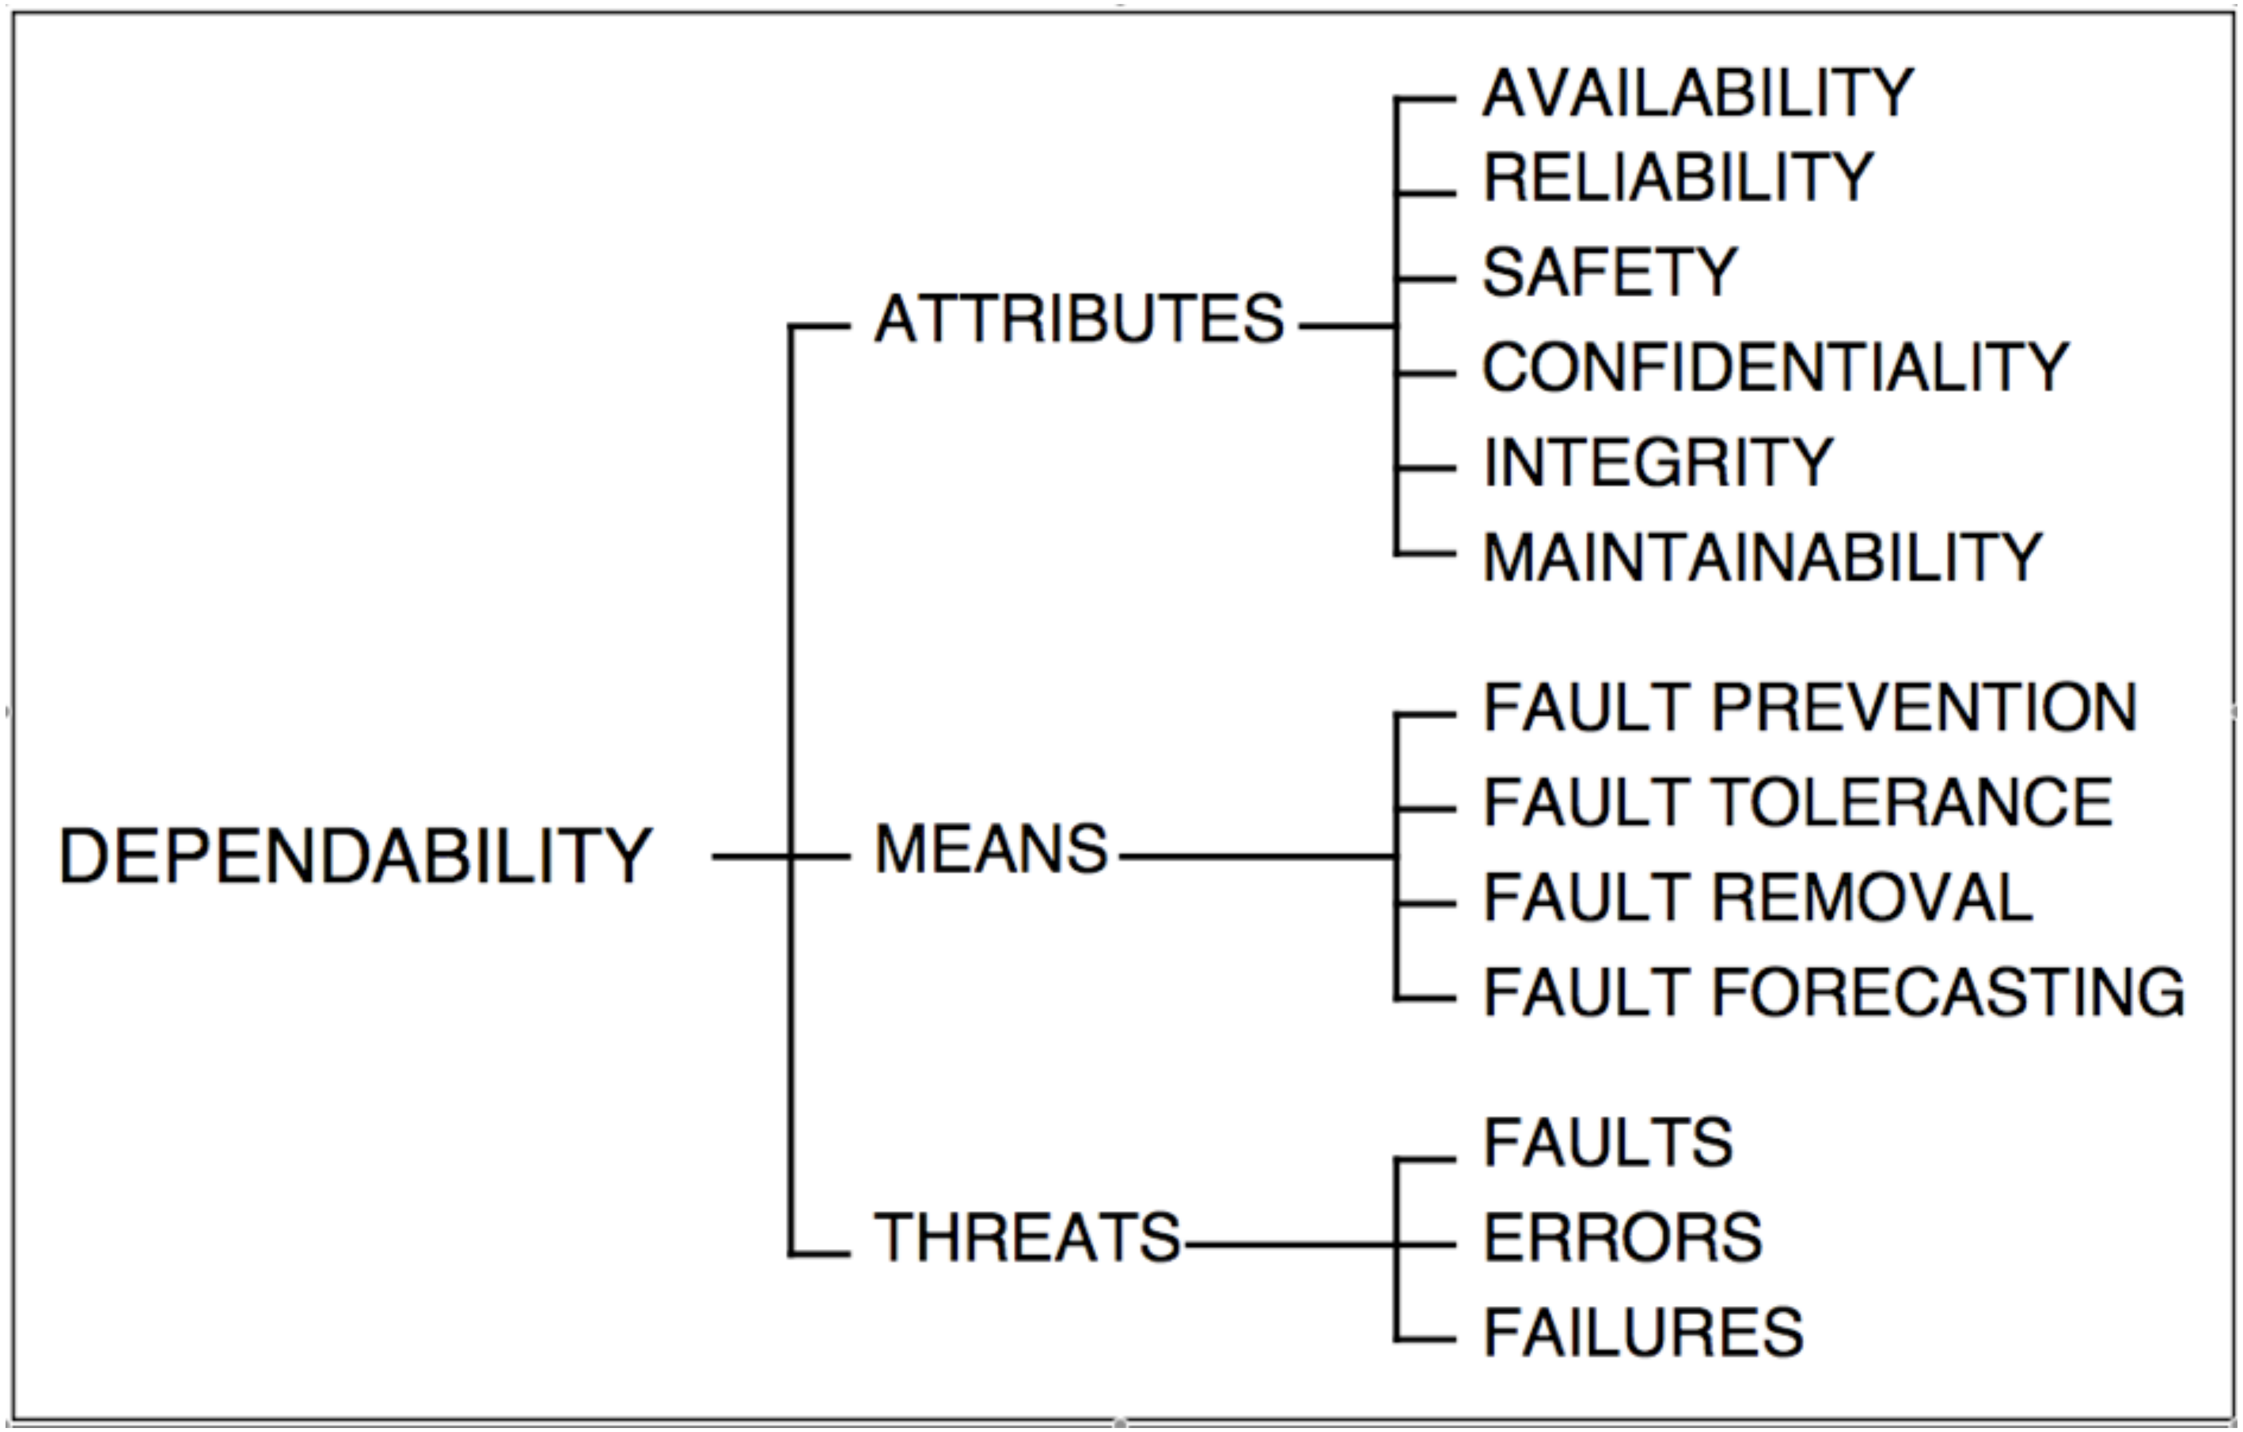

Attributes of Dependability

- **Non-functional attributes** such as reliability and maintainability
- **Complementary nature of viewpoints** in the area of dependability
- In comparison to functional properties
  - ... hard to define
  - ... hard to abstract
  - ... ‚Divide and conquer' does not work as good
  - ... difficult interrelationships
  - ... often probabilistic dependencies

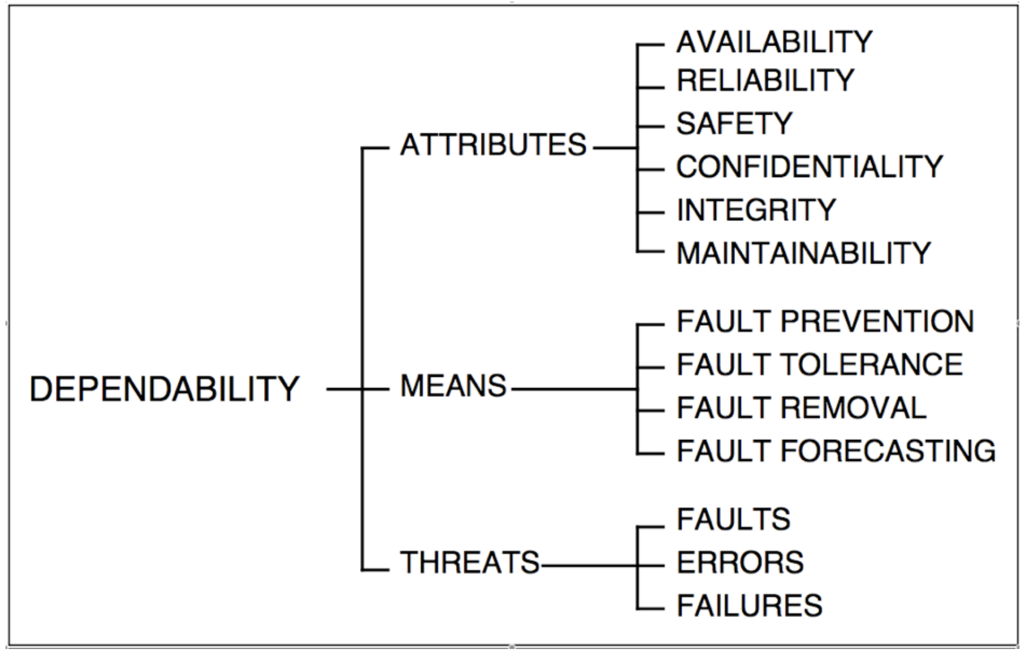

DEPENDABILITY

- ATTRIBUTES
  - AVAILABILITY
  - RELIABILITY
  - SAFETY
  - CONFIDENTIALITY
  - INTEGRITY
  - MAINTAINABILITY
- MEANS
  - FAULT PREVENTION
  - FAULT TOLERANCE
  - FAULT REMOVAL
  - FAULT FORECASTING
- THREATS
  - FAULTS
  - ERRORS
  - FAILURES

Attributes of Dependability

- **Reliability (‚Zuverlässigkeit‘)** - Continuity of service
  - Initial goal for computer system trustworthiness
  - Other disciplines have different understanding
  - *„Reliability is not doing the wrong thing.“* [Gray85]
  - *„Reliability: Ability of a system or component to perform its required functions under stated conditions for a specified period of time*“ [IEEE]
  - *„Reliability is the probability that an item will not fail.“* [Misra]
- **Availability (‚Verfügbarkeit‘)** - Readiness for usage
  - *„Probability that a system is able to deliver correctly its service at any given time.“* [Goloubeva]
  - *„Maintainability is the probability that the item can be successfully restored to operation after failure; and availability ... is a function of reliability and maintainability .“* [Misra]

Observations on Dependability Attributes

- Availability is always required
- Reliability, safety, and security may be optional
- Reliability might be analyzed for hardware / software components
- Availability is always from the system view point

Attributes of Dependability

- **Safety** - Avoidance of catastrophic consequences on the environment
  - Critical applications
  - Specification needs to describe things that should not happen
- **Security** - Prevention of unauthorized access and / or information handling
  - Became especially relevant with distributed systems
- **Confidentiality** - Absence of unauthorized disclosure of information
- **Integrity** - Absence of improper system alteration
  - With respect to either accidental or intentional faults
- **Maintainability** - Ability to undergo modifications and repairs

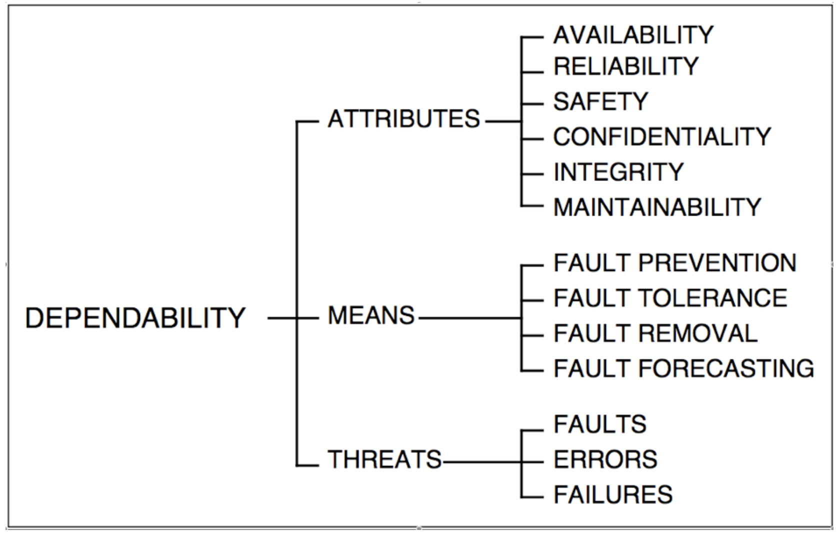

In Detail

- **Reliability** - Function *R(t)*
  - Probability that a system is functioning properly and constantly over time period t
    - Assumes that system was fully operational at t=0
    - Denotes failure-free interval of operation
- **Availability** - Statement if a system is operational at a point in time / fraction of time
  - Describe system behavior in presence of error treatment mechanisms
  - **Instantaneous availability (at t)** - Probability that a system is performing correctly at time t; equal to reliability for non-repairable systems: $A_i(t) = R(t)$
  - **Steady-state availability** - Probability that a system will be operational at any random point of time, expressed as the fraction of time a system is operational during its expected lifetime: $A_s = Uptime / Lifetime$

Probability of Events

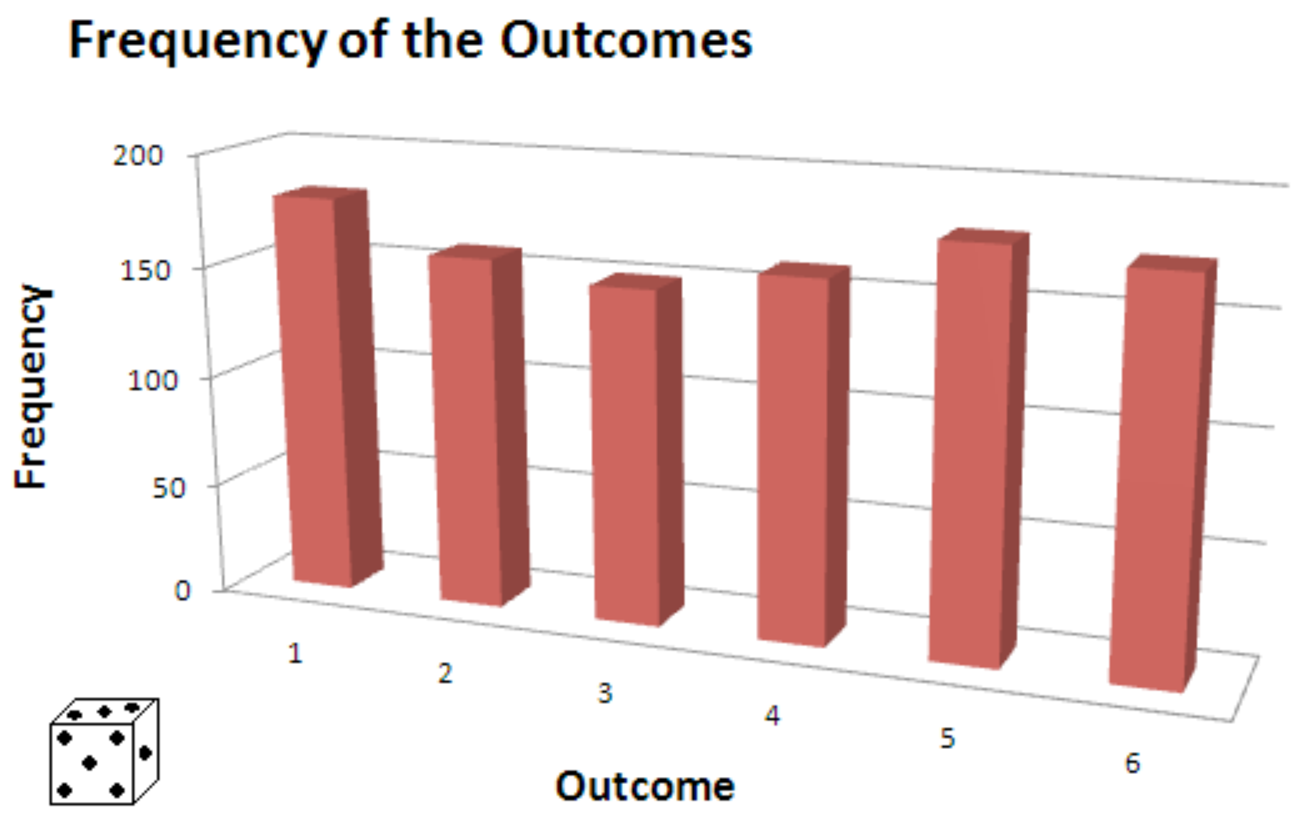

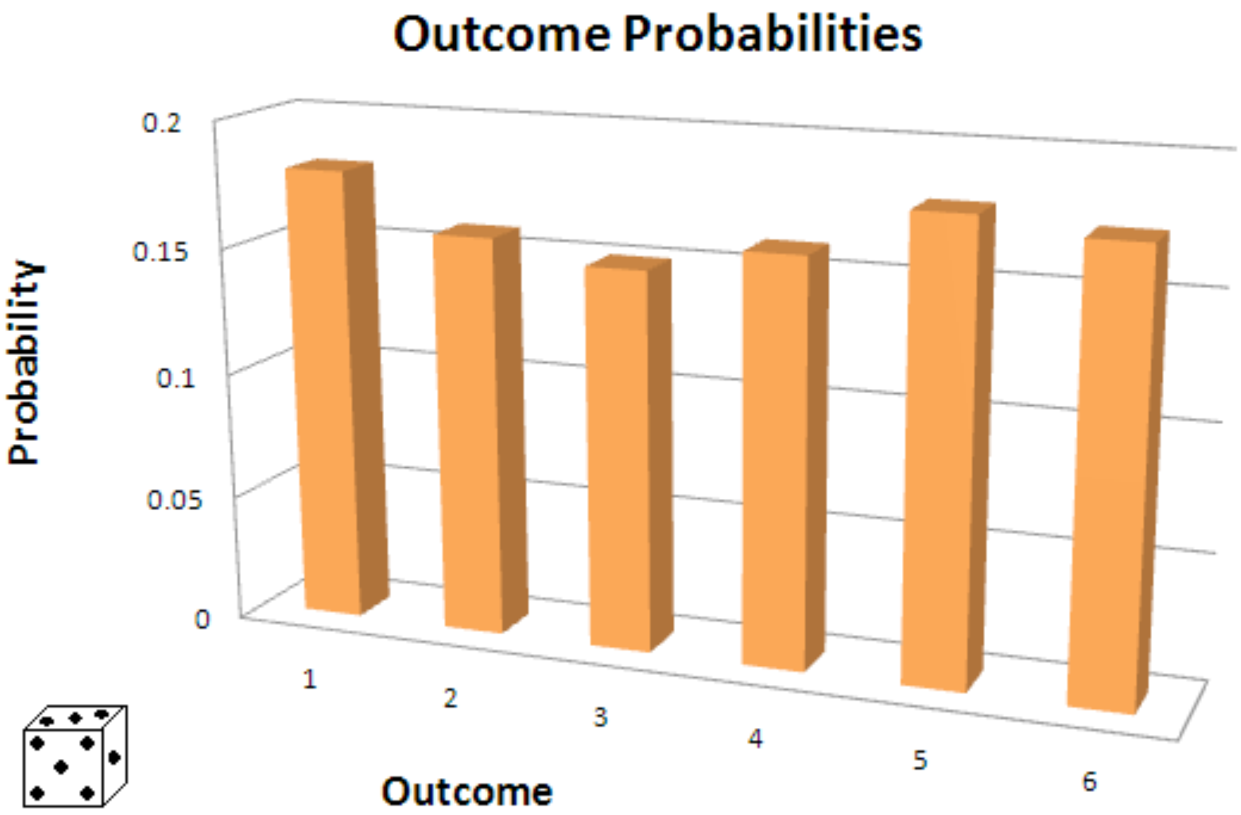

PDF & CDF

- Probability density function *pdf* for random variable *X*
  - Discrete random variable: Probability that *X* will be *x*
  - Continuous variable: Probability that *X* is in *[a, b]*
    - Computed as area under the density function in this range

$$P(a \leq X \leq b) = \int_a^b f(x)dx \text{ and } f(x) \geq 0 \text{ for all } x$$

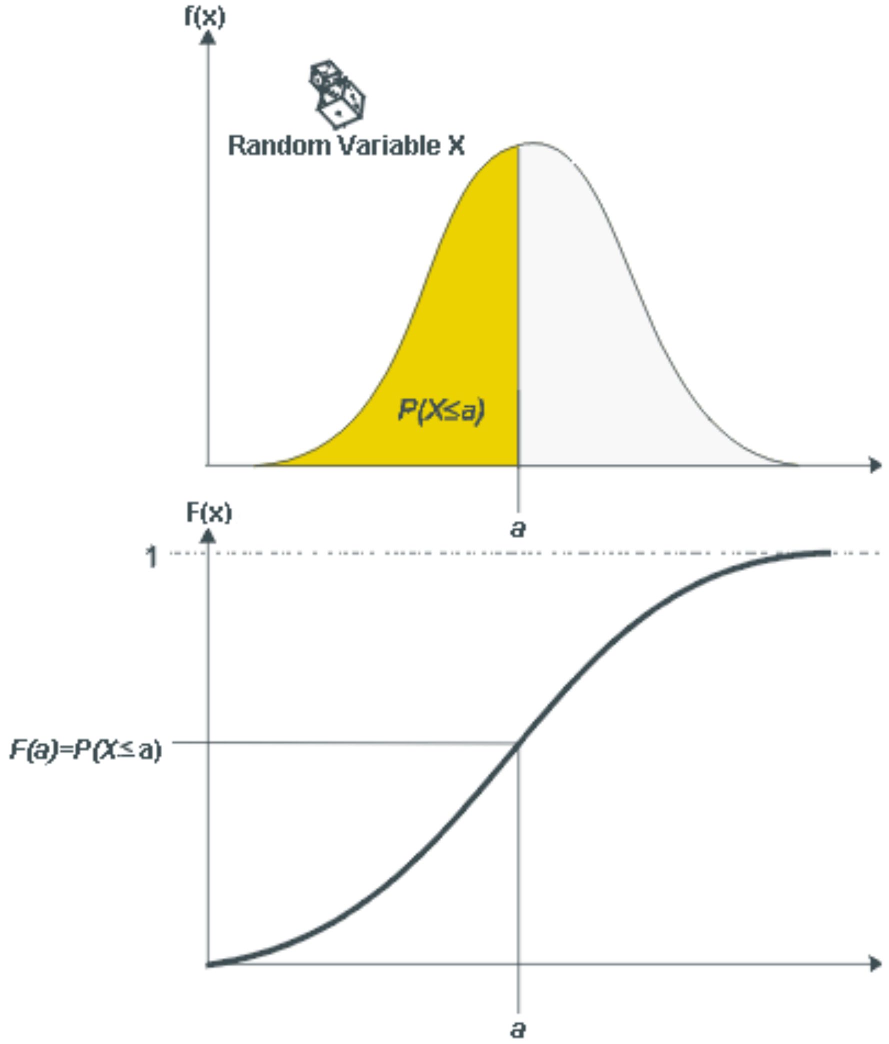

(C) weibull.com

- Cumulative distribution function *cdf(x)*: Probability that the value of the random variable is at most x

$$F(x) = P(X \leq x) = \int_{0,-\infty}^{x} f(s)ds$$

  - Limits of integration depend on the nature of the distribution function
- Value of *cdf* at *x* is always area under *pdf* from 0 to *x*

PDF Examples

- Well-known statistical distributions, each describing a random variable behavior
- Continuous version described by PDF (discrete pendant would be histogram)

Normal distribution
(mean, variance)

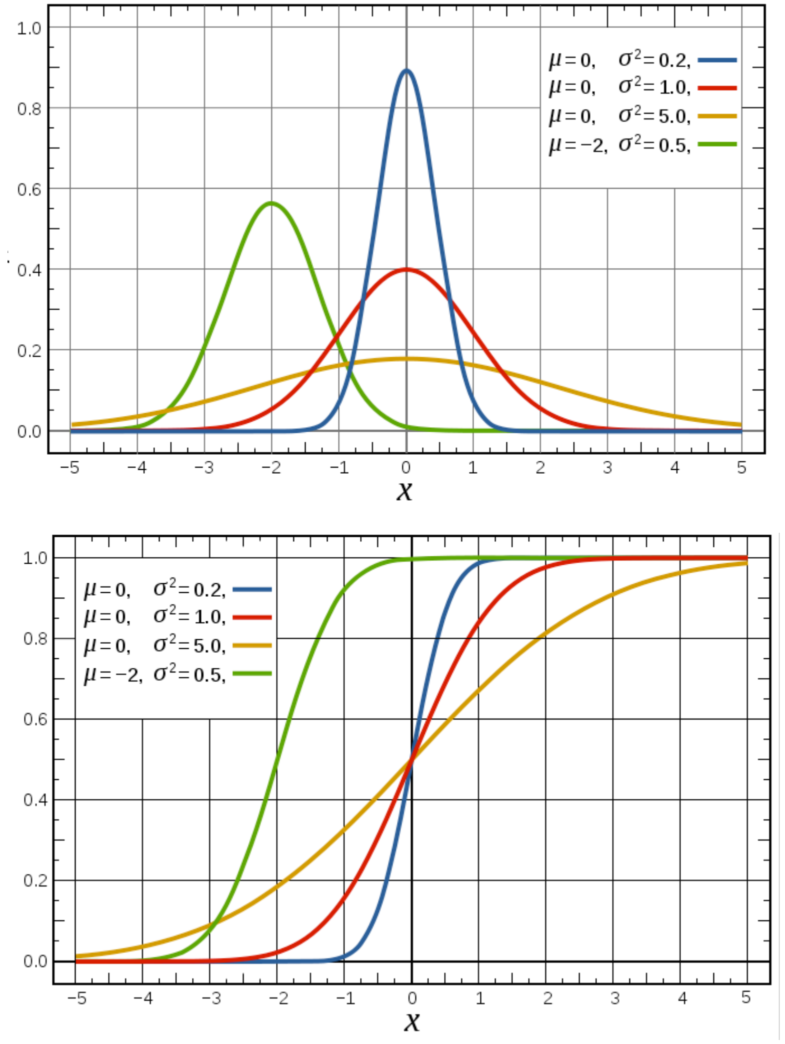

Exponential distribution
(rate parameter)

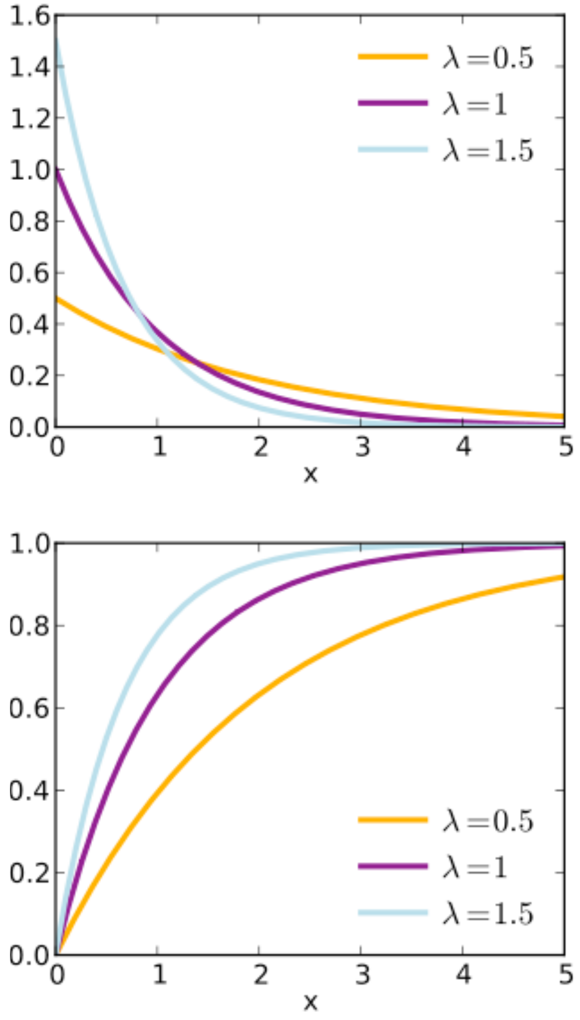

Probability density function

$$f_\lambda(x) = \begin{cases} \lambda e^{-\lambda x} & x \geq 0 \\ 0 & x < 0 \end{cases}$$

Cumulative distribution function

$$F(x) = \int_{-\infty}^{x} f_\lambda(t)\ \mathrm{d}t = \begin{cases} 1 - e^{-\lambda x} & x \geq 0, \\ 0 & x < 0. \end{cases}$$

The Reliability Function R(t)

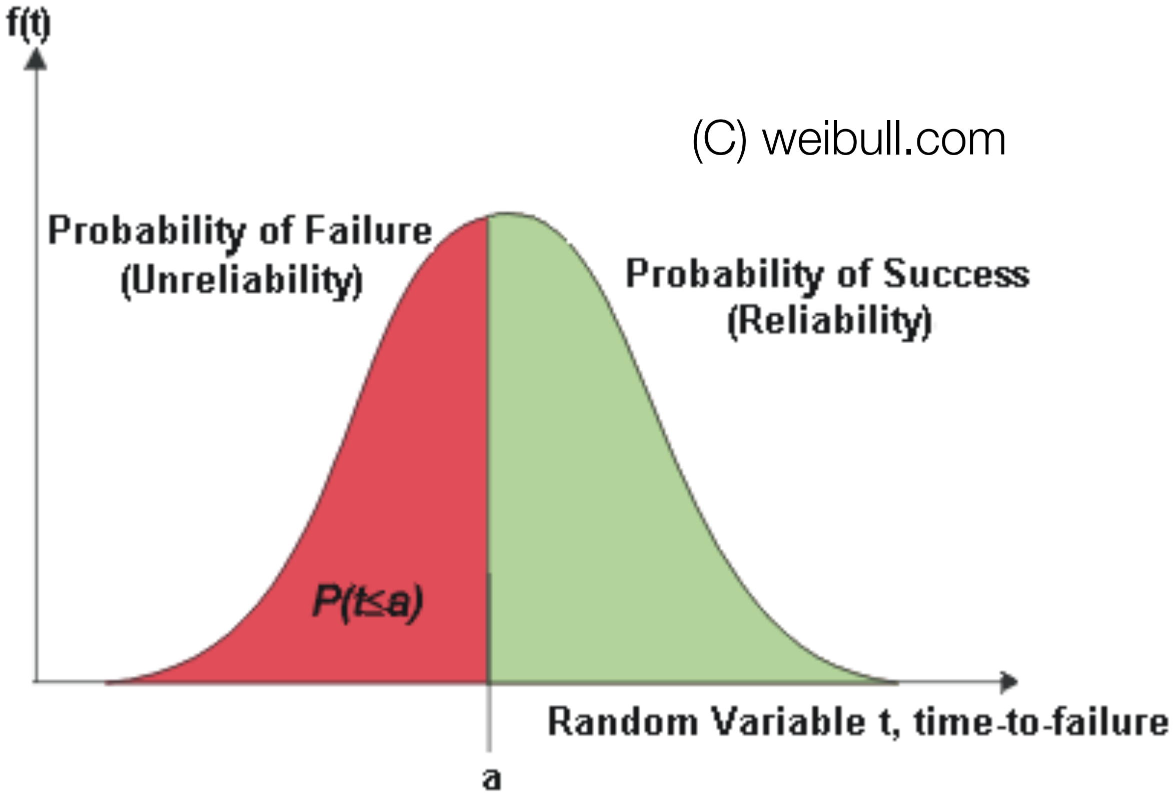

(C) weibull.com

- Reliability: Probability *R(t)* that a component works for time period [0,*t*]
- Idea: Express time period of correct operation as **continuos random variable *X***
  -> **time to failure**
  - *cdf(t) of this variable:* Describes probability of failure before t -> **Unreliability Function F(t)**
  - *1-cdf(t):* Describes probability of a failure after t -> time to failure -> **Reliability Function R(t)**
    - This works since (A) working / non-working is a binary decision, (B) the area under the complete *pdf* is 1 and (C) the ‚red‘ area is the result of the *cdf* function
- Typically, failures are modeled as Poisson process
  - Poisson properties lead to exponential distribution for the time between events
  - This time therefore only depends on failure rate parameter

Failure Rate

- Time to failure is not always measured in calendar time, might be discrete
  - Number of kilometers driven with a car
  - Number of rotations / cycles for a mechanical component
- Treat *pdf* for time-to-failure random variable X as **failure density function**
  - Can be computed as derivative of the unreliability function

$$f(t) = dF(t)/dt$$

- **Failure rate** / hazard rate function - mean frequency of failures at time t
  - Conditional probability of a failure between a and b, given the survival until t

$$\lambda(t) = \frac{f(t)}{R(t)} = \lambda \text{ for constant failure rate}$$

Why Exponential ?

- Distribution function that models the **memoryless property** of the Poisson process
    - $P(T > t + s|T > t) = P(T > s)$, e.g. $P_{Failure}(5 \text{ years}|T > 2 \text{ years}) = P_{Failure}(3 \text{ years})$
    - Failure is not the result of wear-out
    - Models ‚intrinsic failure‘ behavior, assumed for the majority of hardware life time
    - Weibull distribution as alternative, can also model **tear-in** and **wear-out**
- Some natural phenomena have constant failure rate (e.g. cosmic ray particles)

The Reliability Function R(t)

- Failures occur continuously and independently at a constant average rate (Poisson process)
- Increasing probability of failure with increasing t - *cdt* function
- Failure rate λ from experience or complexity measures
- Cumulative distribution function:

$$F(x;\lambda) = \begin{cases} 1 - e^{-\lambda x}, & x \geq 0, \\ 0, & x < 0. \end{cases}$$

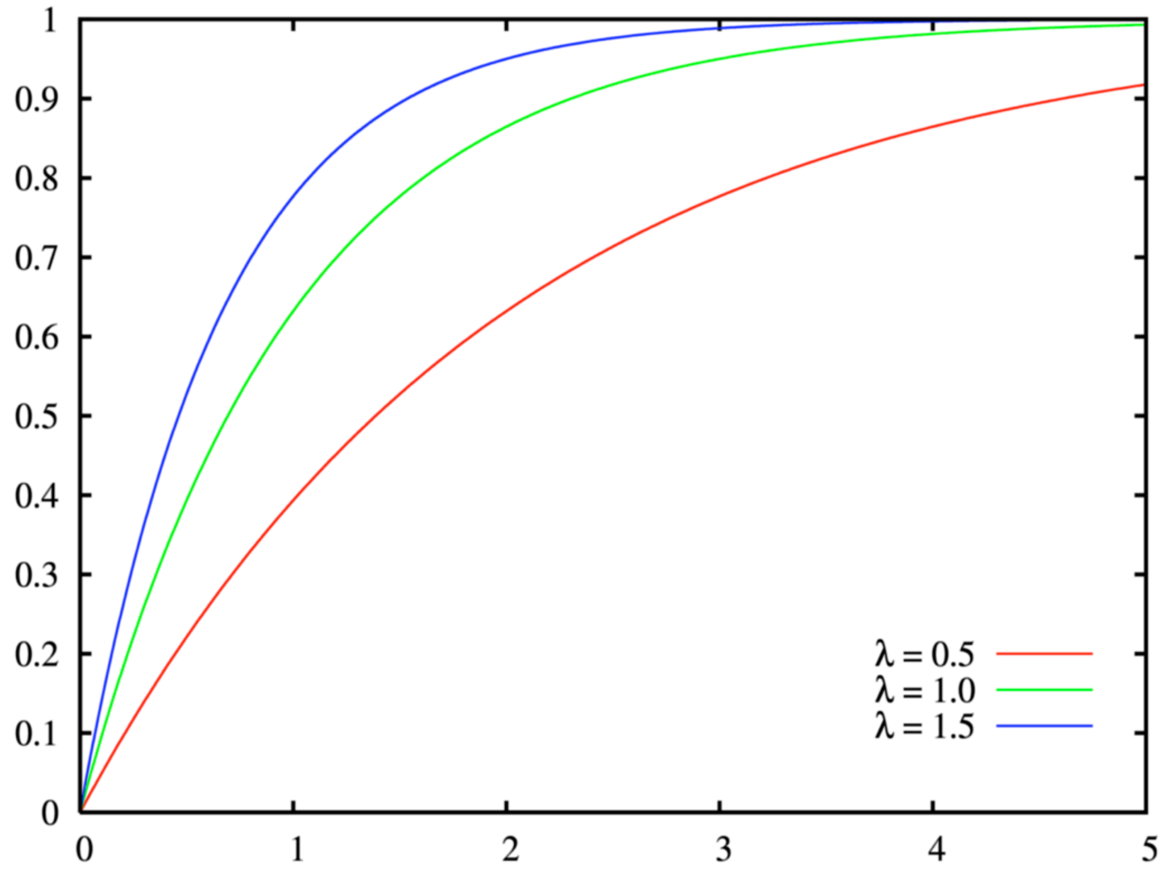

CDF: Probability of a failure before t

- Reliability function (survival probability) for exponential failure distribution:

$$R(t) = P(X > t) = 1 - F(t) = e^{-\lambda x} \text{ with } F(x) = 1 - e^{-\lambda x}$$

Variable Failure Rate in Real World

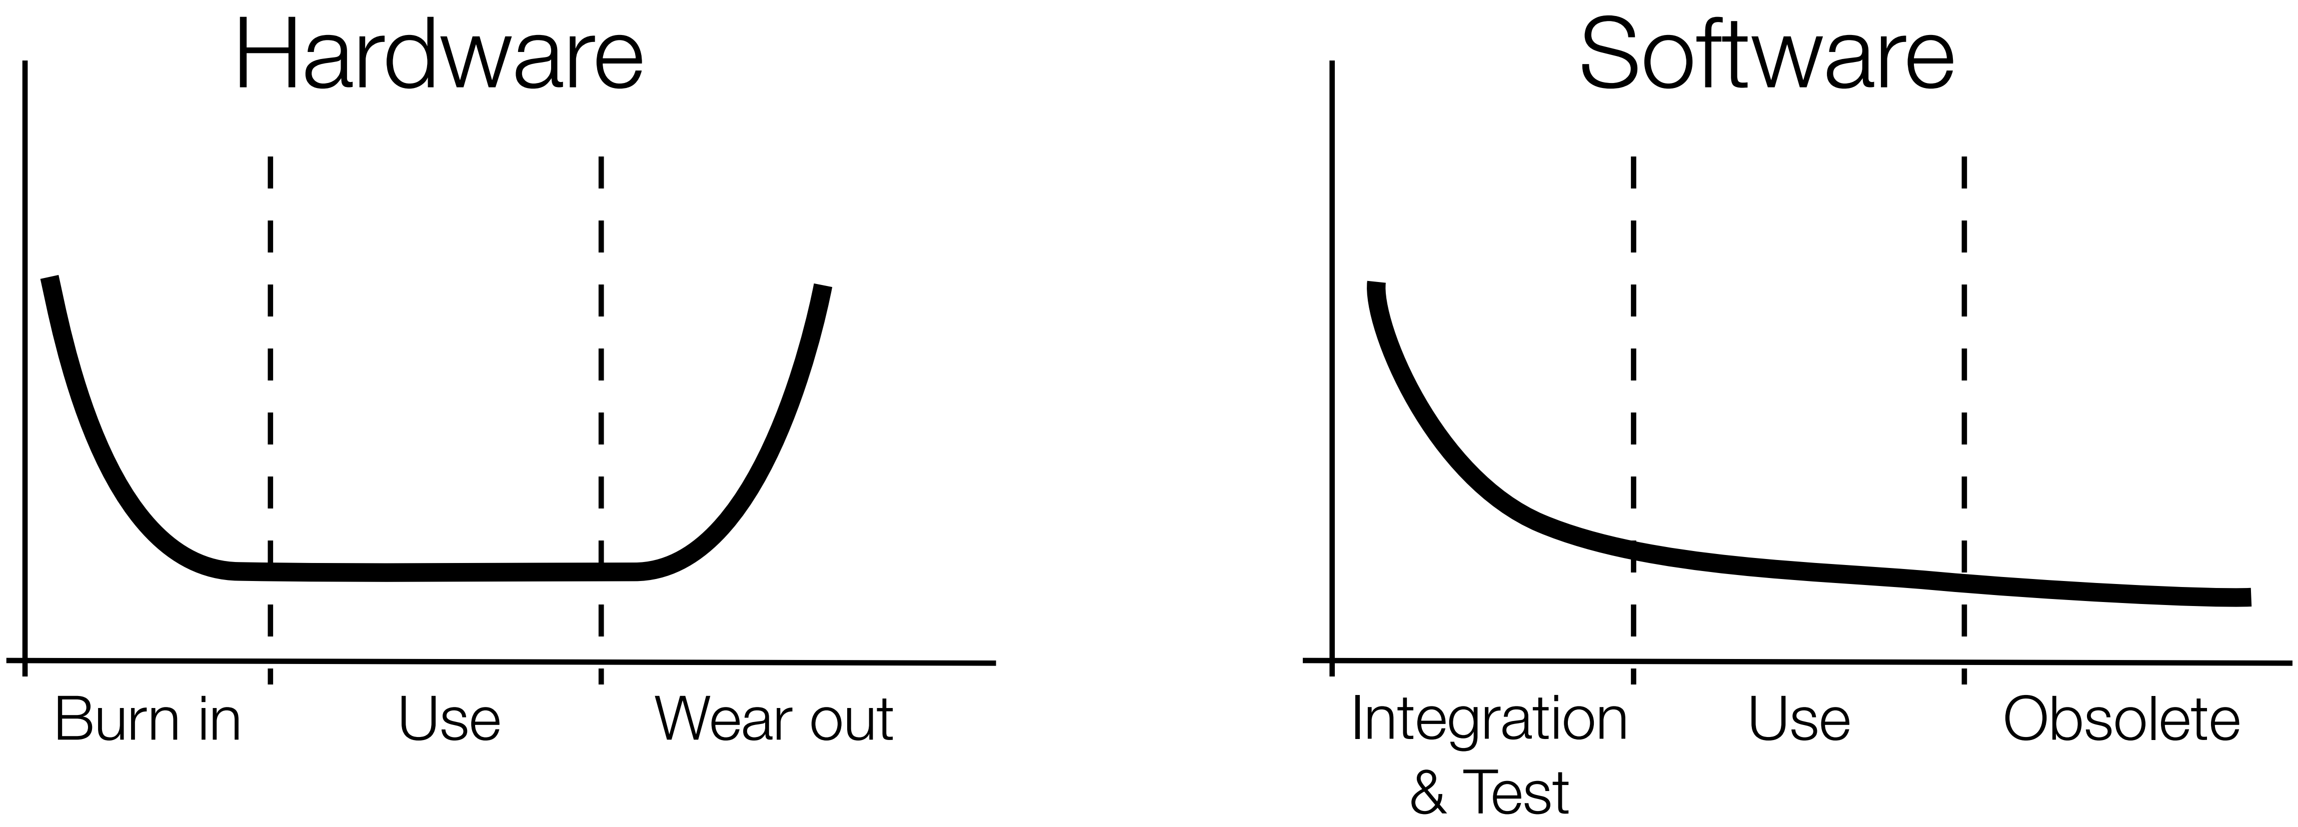

- Failure rate is treated as constant parameter of the exponential distribution
- (maybe invalid) simplification, mostly combined solution in practice:
    - Exponential distribution when failure rate is constant
    - Weibull distribution when failure rate is monotonic decreasing / increasing

Weibull Distribution

- Most widely used life distribution beside exponential
- Named after Swedish professor Waloddi Weibull (1887 - 1979)
  - Originally invented for modeling material strength
- Very flexible through parametrization, can model many other probability distribution types
  - k: Shape parameter
    - k=1: constant hazard rate, like exponential distribution
    - k<1: Hazard rate decreases over time (‚infant mortality‘)
    - k>1: Hazard rate increases with time (‚aging‘), like with (log)normal distribution
  - λ: Scale parameter

$$f(x;\lambda,k) = \begin{cases} \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k} & x \geq 0, \\ 0 & x < 0, \end{cases}$$

pdf

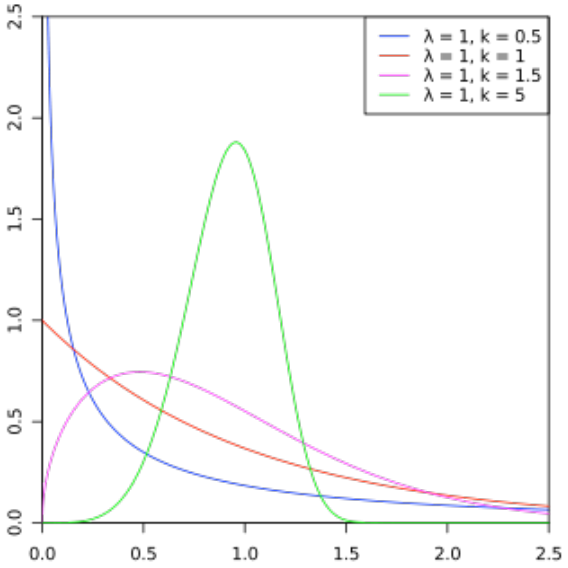

cdf

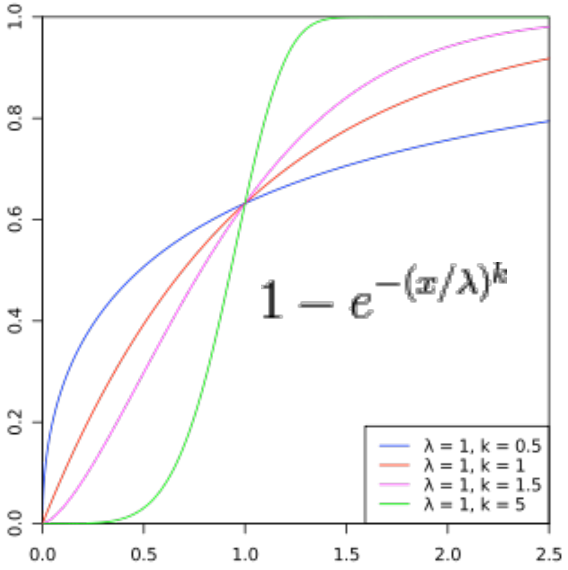

Hardware Failure Rate

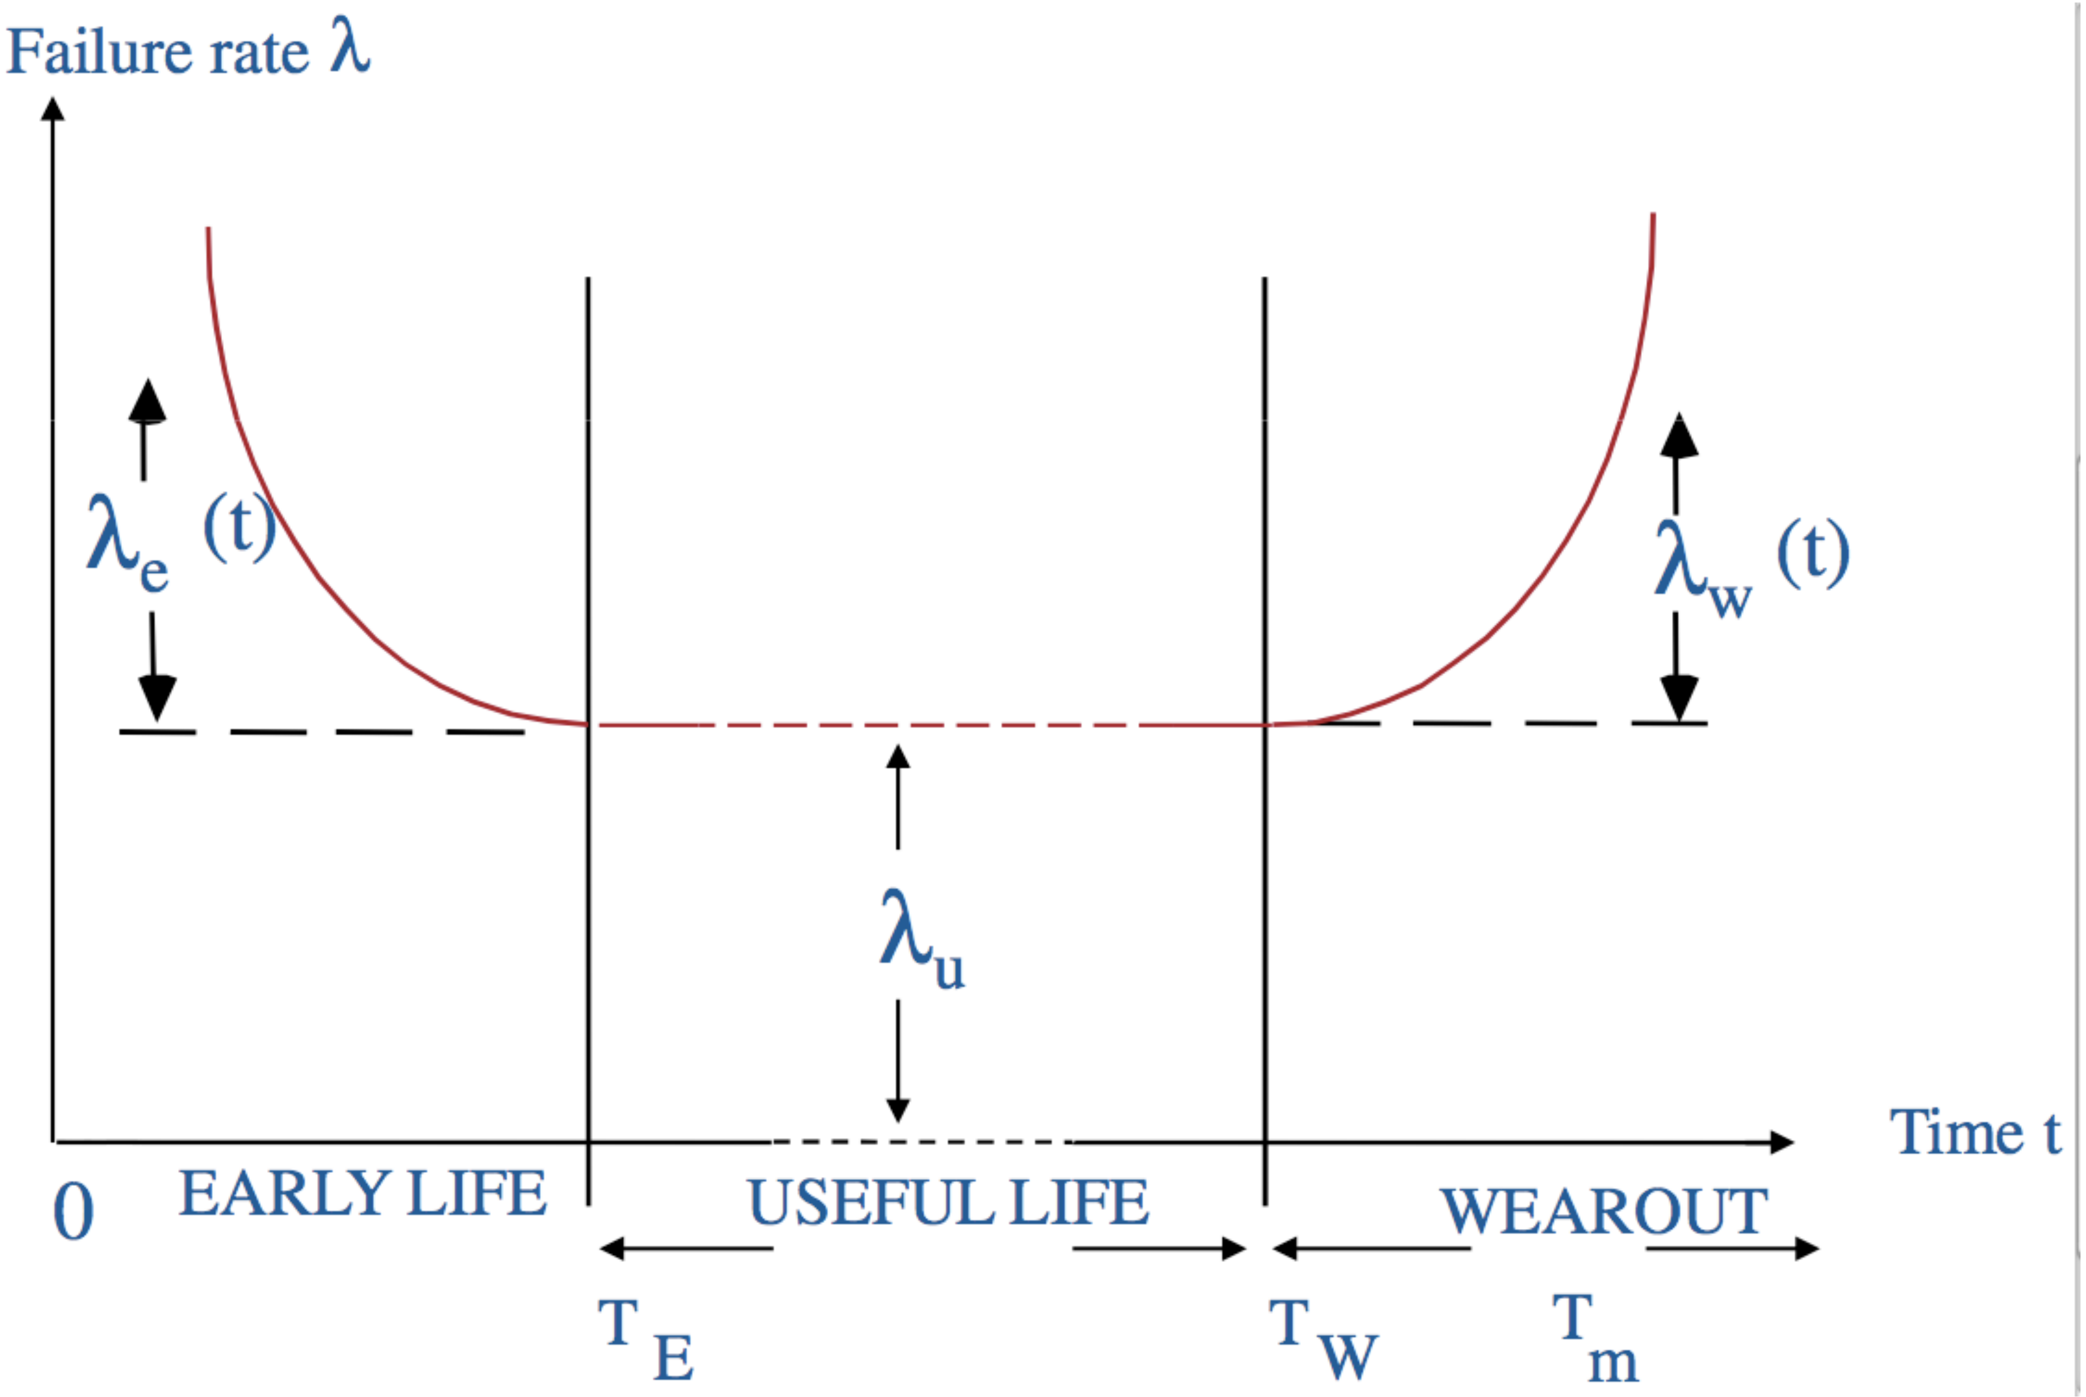

Software Failure Rate

- Industrial practice
- When do you stop testing ?   -   No more time, or no more money ...

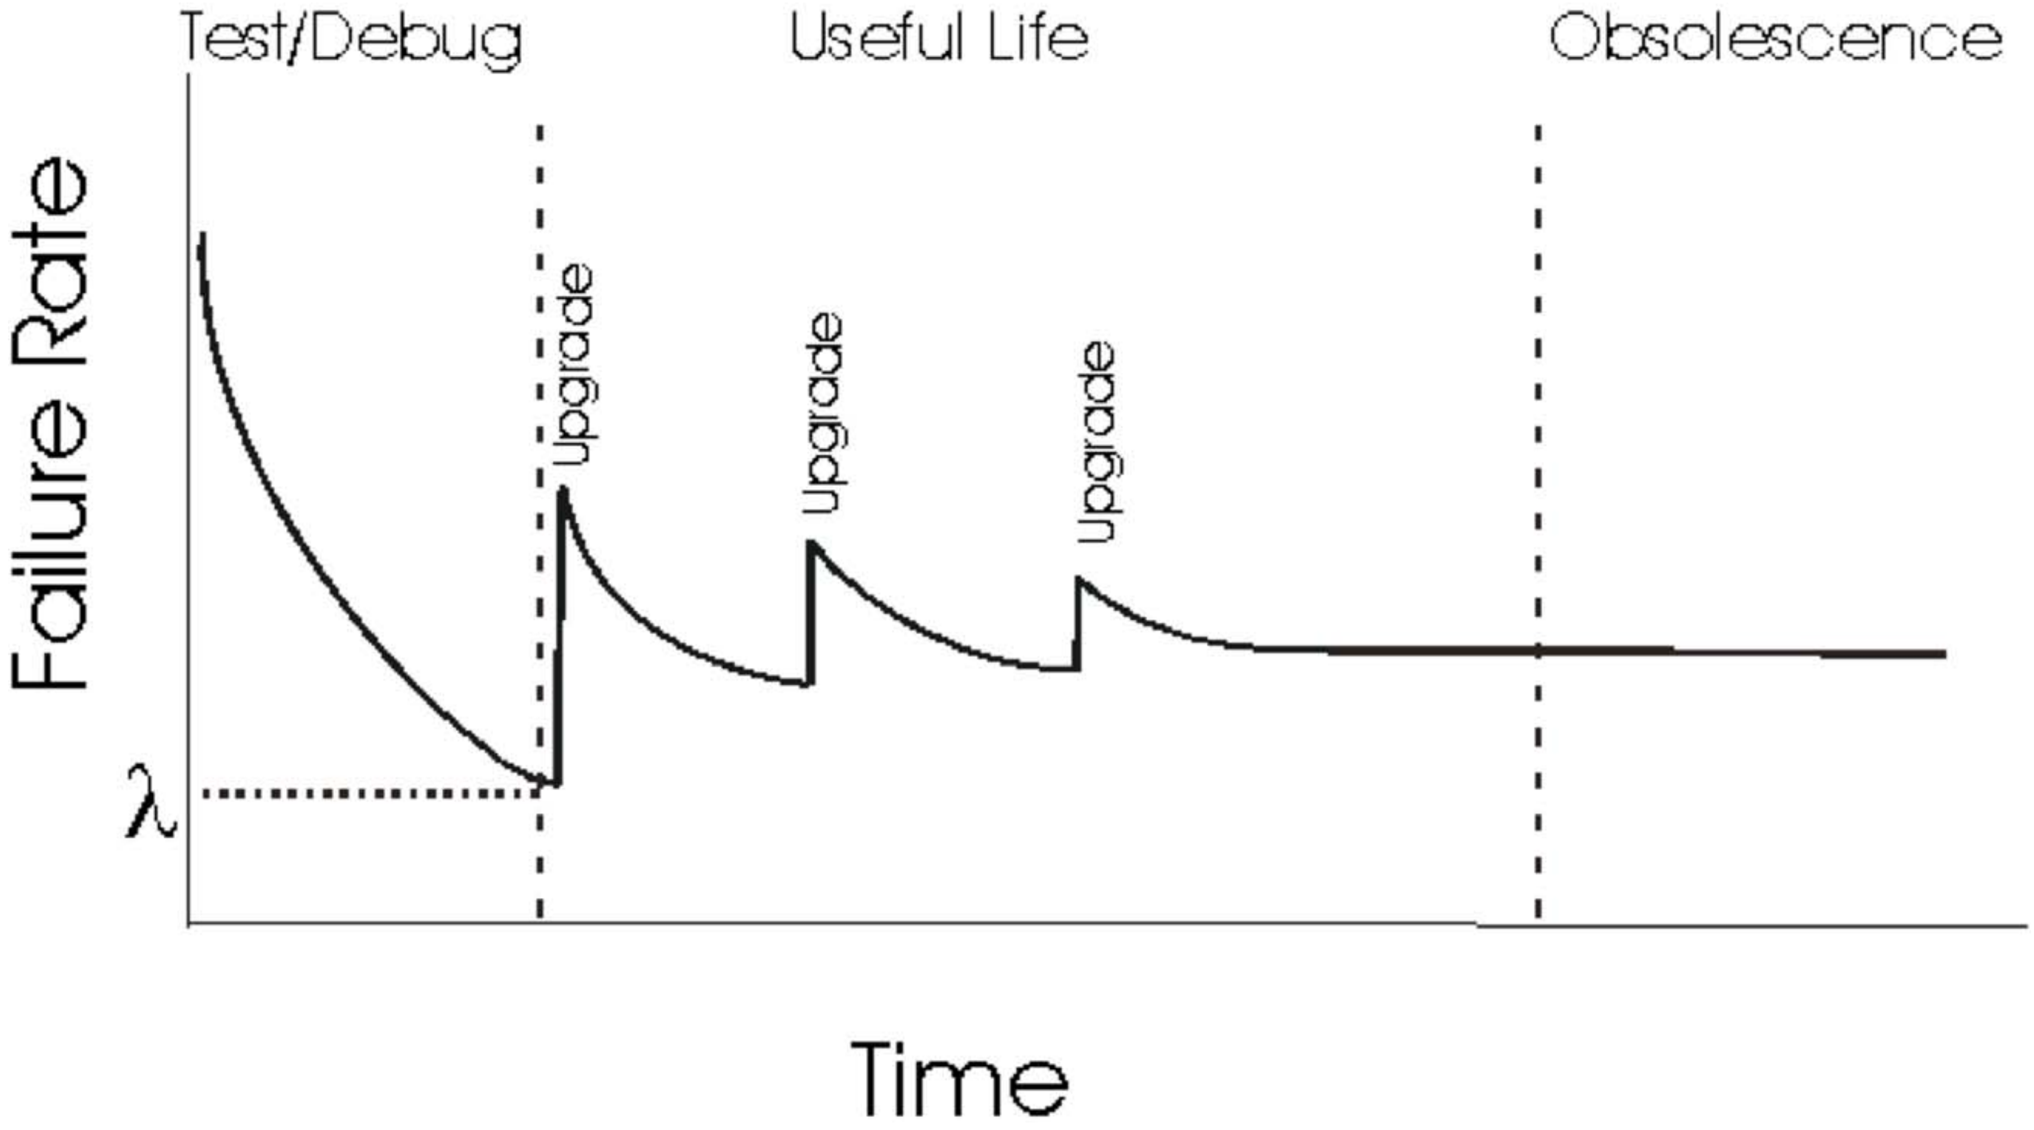

(C) Malek

Failure Rate Examples

- Standards from experience provide base data for component reliability
- Society of Automotive Engineers (SAE) reliability model

$$\lambda_p = \lambda_b \Pi_{i=1}^{b} \pi_i$$

- Predicted failure rate λ_p
- Base failure rate for the component λ_b
- Various modification factors π_i
  - Component composition
  - Ambient temperature
  - Location in the vehicle

Life-Stress Relationship

- Formulate a model that includes the life distribution and how outside factors (such as stress) change this distribution
- Example - **load sharing redundancy**: Component reliability depends indirectly on the number of previously failed components
  - Single component failure distribution is no longer sufficient for describing reliability
  - Model must describe the effect of load and the failure probability
- Typical approach: Define load-dependent parameter of the distribution function

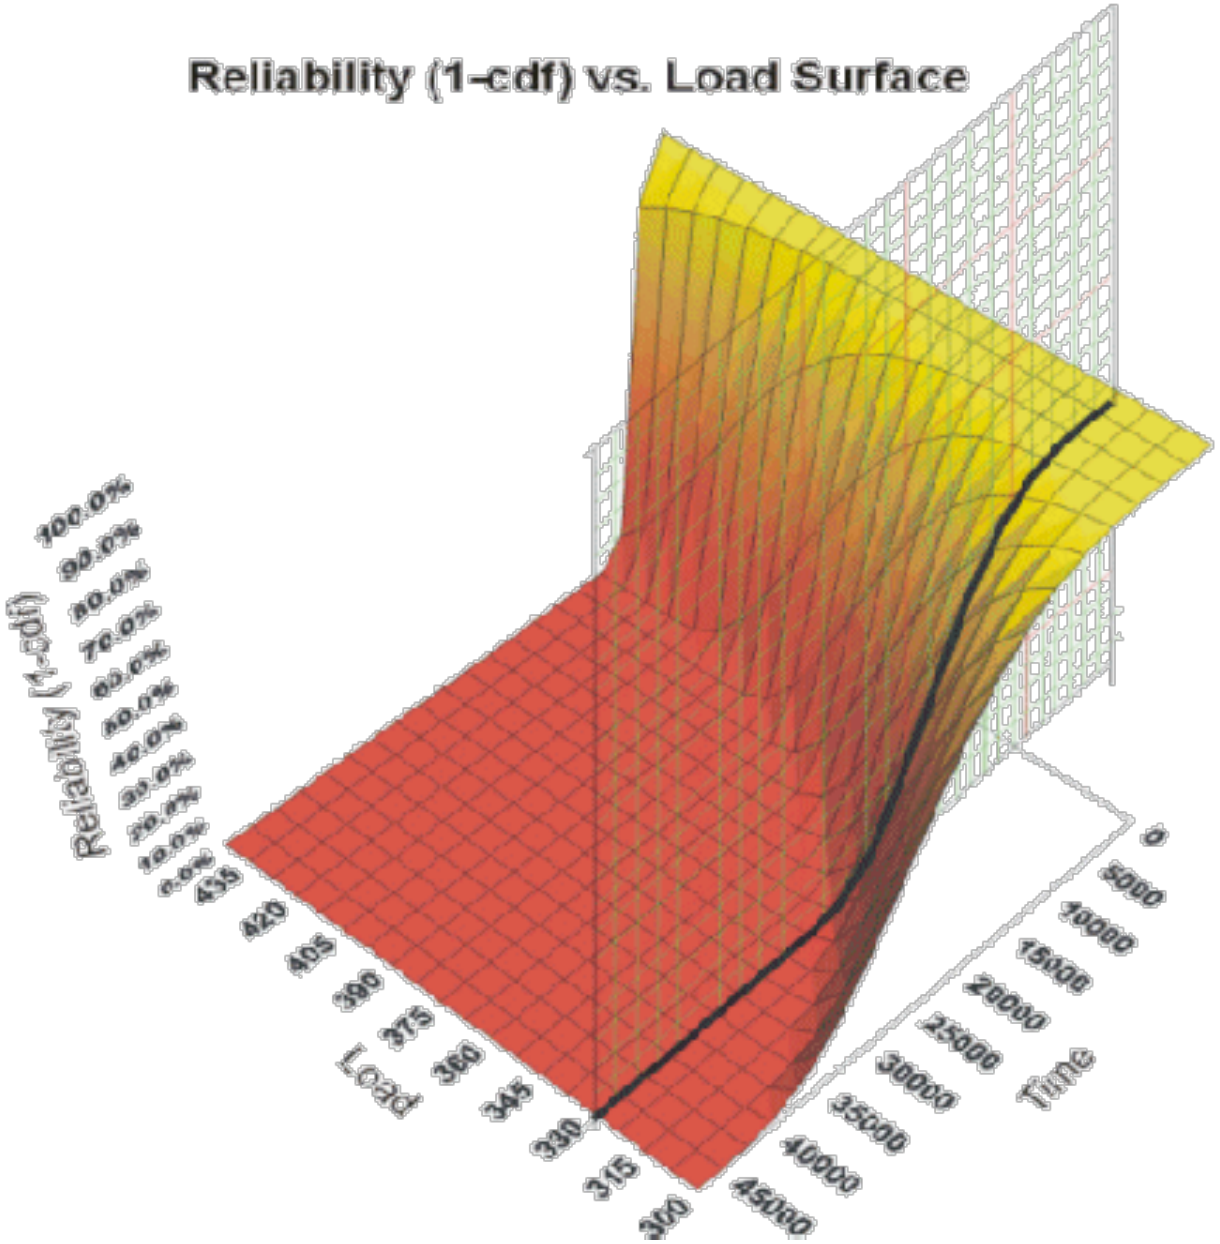

(C) weibull.com

Other Reliability Distributions

- (Mixed) Weibull distribution
- Normal distribution
- Lognormal distribution
- (Generalized) Gamma distribution
- Logistic distribution
- Loglogistic distribution

Steady-State Availability

- **Mean time to failure (MTTF)** -
  Average time it takes to fail
  -> average uptime

- **Mean time to recover / repair (MTTR)** -
  Average time it takes to recover

- **Mean time between failures (MTBF)** -
  Average time between two successive failures

- **Availability** = Uptime / Lifetime
  = MTTF / MTBF

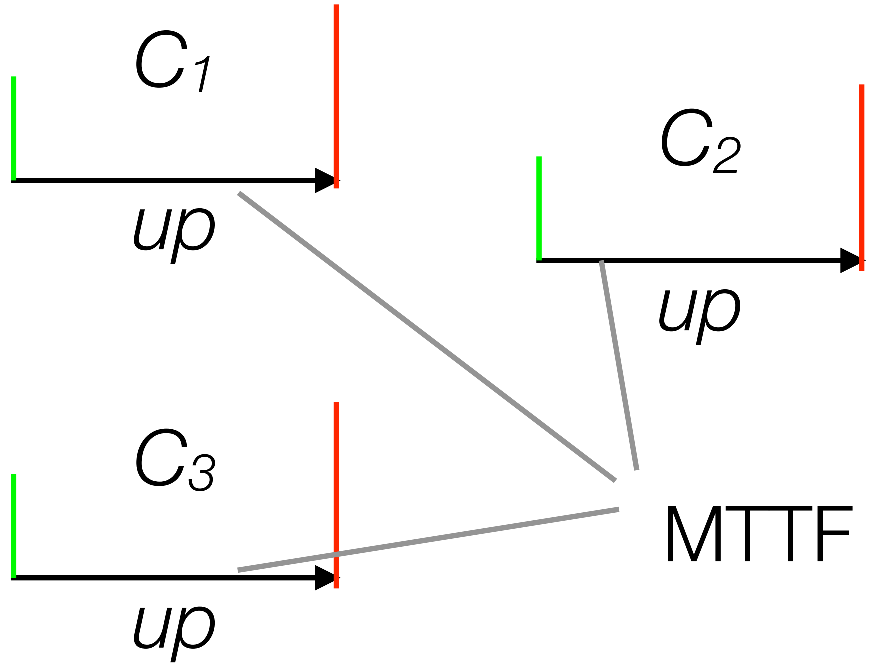

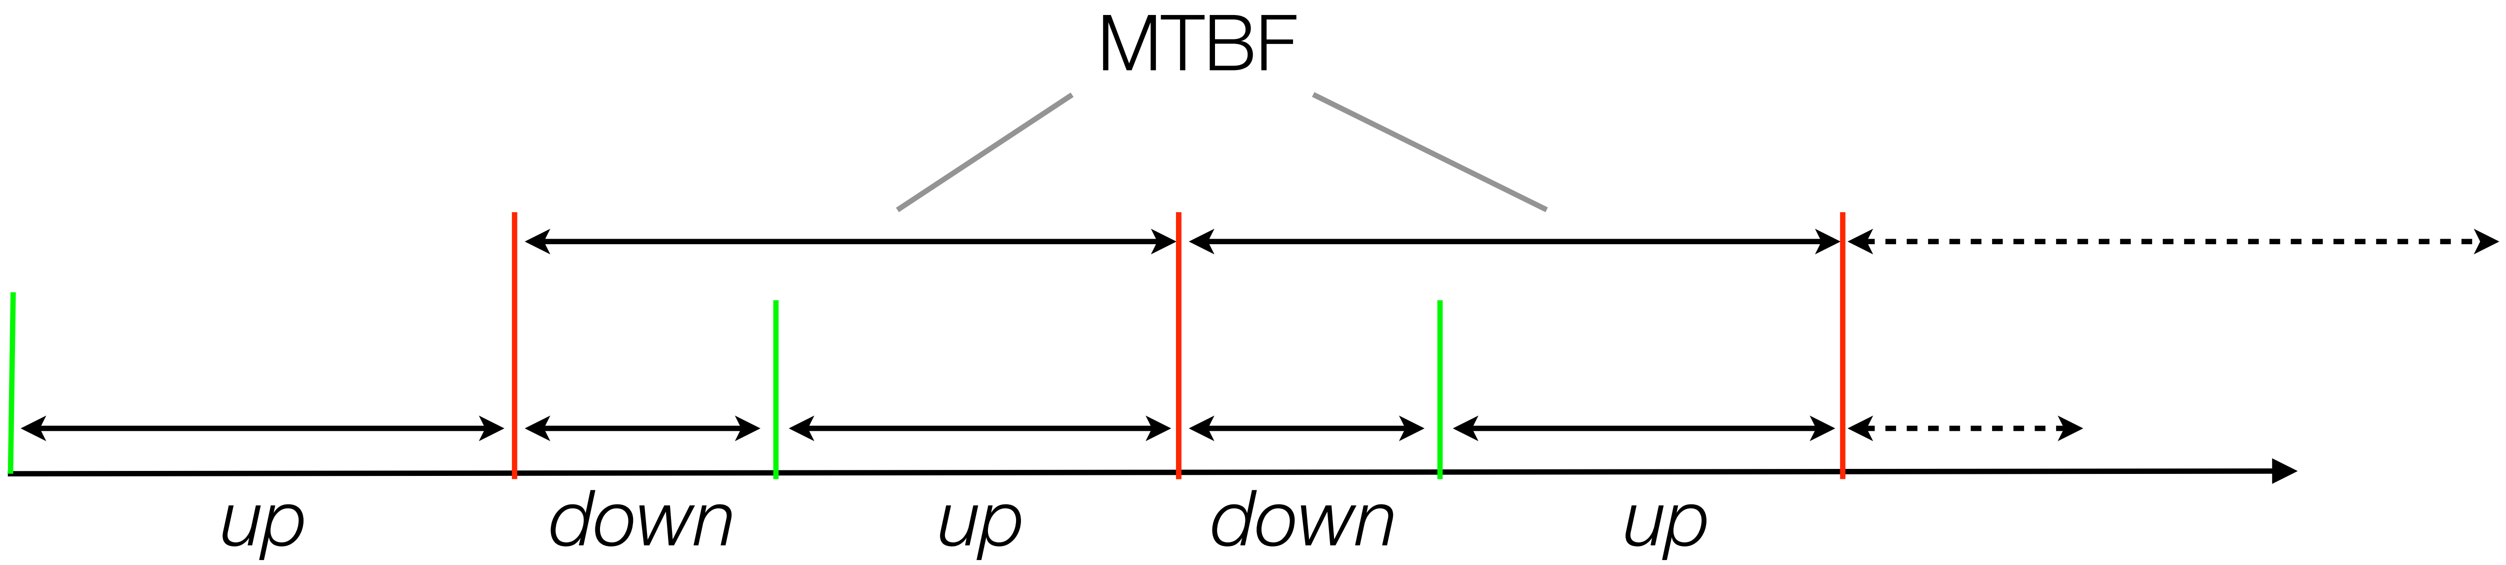

Steady-State Availability and MTBF

- Expressing availability with MTTF/MTBF describes a repairable systems behavior under infinite time assumption
  - MTTF and MTTR get stable over longer time periods
  - Fulfils also the steady-state condition
- Expressing *dependability* with MTTF ‚should‘ imply a non-repairable system, expressing *dependability* with MTBF ‚should‘ imply a repairable system
- Sometimes **MTBF** means **mean time BEFORE failure** = MTTF
  -> typical source of confusion
- Exponential distribution:
  Reciprocal of the rate parameter is equivalent to the *distribution mean*
  (Example: With 4 events per hour, you can expect one roughly every 15 minutes).

$$\lambda = \frac{1}{MTTF}$$

Example

- Test population with 50 HDDs and 100 hours of testing, 2 drives fail during the test
  - As usual, we assume exponential distribution of the time to failure
  - Reliability at t=100 is known to be 96% (48/50)
  - Reciprocal of the according failure rate is the MTTF

$$R(t) = P(X > t) = 1 - F(t) = e^{-\lambda x} \text{ with } F(x) = 1 - e^{-\lambda x}$$

$$R(100hours) = e^{-\lambda 100} = 0.96$$

$$\lambda = -\frac{ln 0.96}{100} = 0,000408$$

$$MTTF = \frac{1}{\lambda} = 2449,66hours$$

MTBF / MTTF in Practice

- Often express average failure behavior (statistics) for a component population
- Good for relative comparison, not for expected life time expectation of one unit
- Example: Hard disk with MTTF of 500.000 hours and 5 years of expected operation (‚service life‘)
  - Drive of this type is expected to run 5 years without problems
  - Large group of such drives will (on average) have one failed drive after 500.000 hours of **accumulated** life time
- MTBF in practice is a weak approximation - sum of up-phases / number of failures
  - Assumes renewable system with homogeneous failure distribution over lifetime

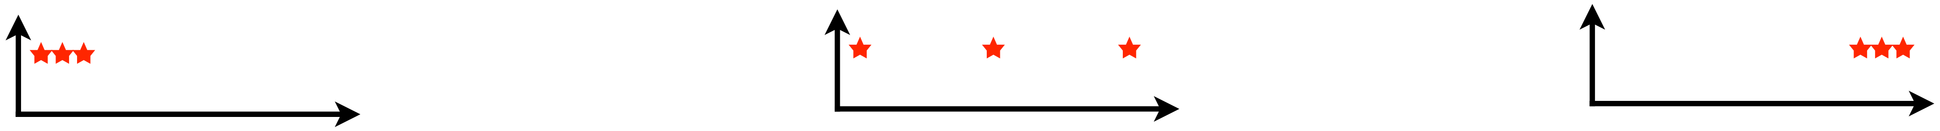

Steady-State Availability

$$A = \frac{Uptime}{Uptime+Downtime} = \frac{MTTF}{MTTF+MTTR}$$

Availability	*Downtime per year*	*Downtime per week*
90.0 % (1 nine)	*36.5 days*	*16.8 hours*
99.0 % (2 nines)	*3.65 days*	*1.68 hours*
99.9 % (3 nines)	*8.76 hours*	*10.1 min*
99.99 % (4 nines)	*52.6 min*	*1.01 min*
99.999 % (5 nines)	*5.26 min*	*6.05 s*
99.9999 % (6 nines)	*31.5 s*	*0.605 s*
99.99999 % (7 nines)	*0.3 s*	*6 ms*

Amazon EC2 SLA (2012)

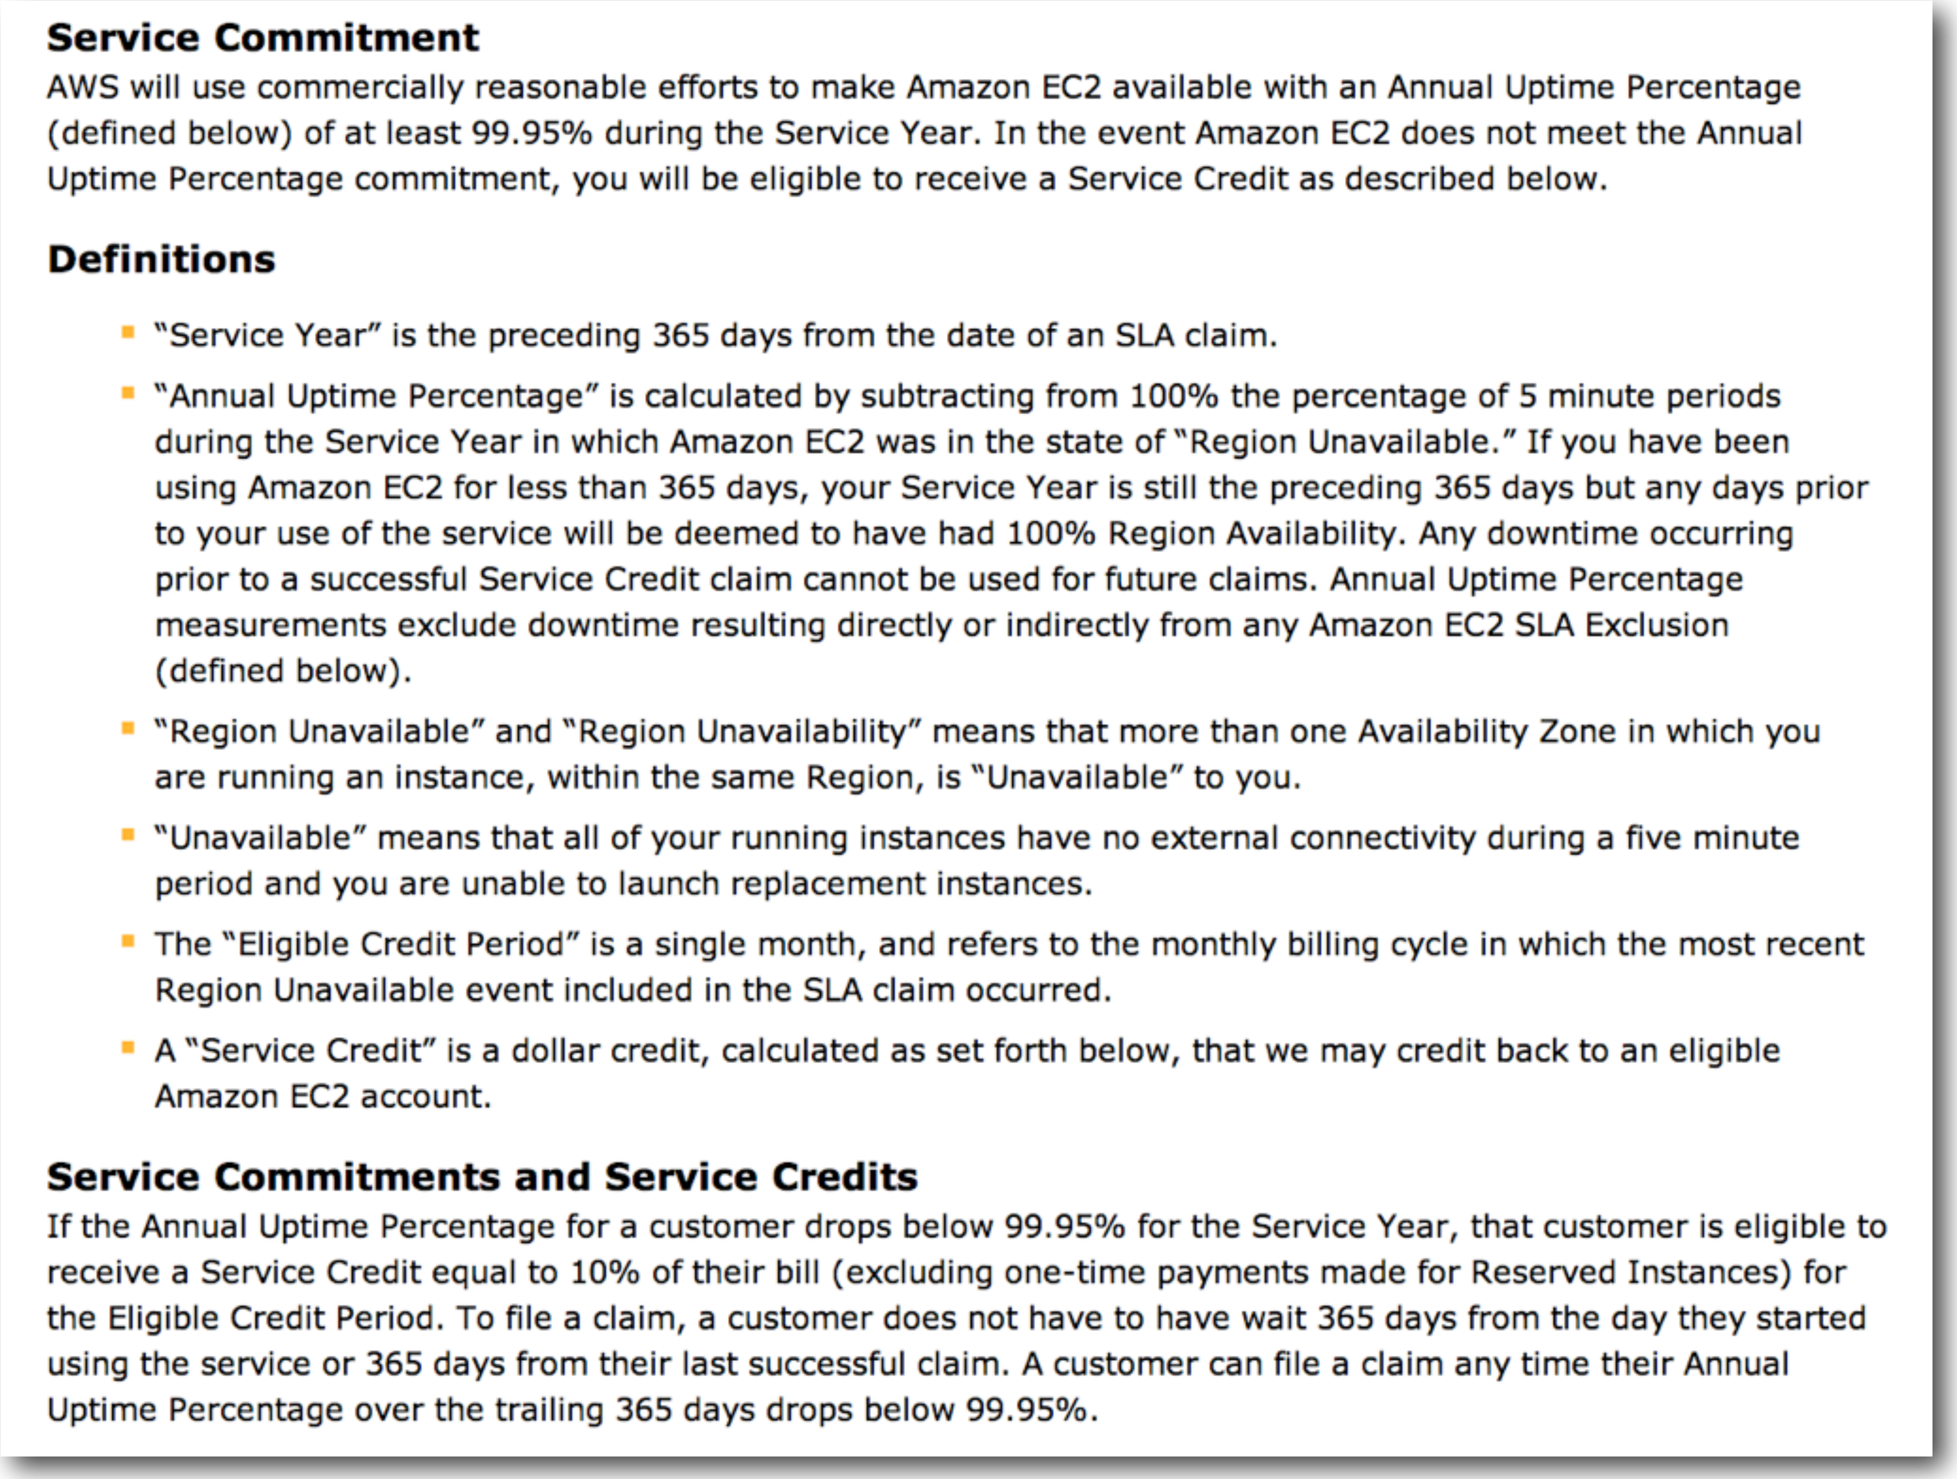

Service Commitment

AWS will use commercially reasonable efforts to make Amazon EC2 available with an Annual Uptime Percentage (defined below) of at least 99.95% during the Service Year. In the event Amazon EC2 does not meet the Annual Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Service Year" is the preceding 365 days from the date of an SLA claim.
- "Annual Uptime Percentage" is calculated by subtracting from 100% the percentage of 5 minute periods during the Service Year in which Amazon EC2 was in the state of "Region Unavailable." If you have been using Amazon EC2 for less than 365 days, your Service Year is still the preceding 365 days but any days prior to your use of the service will be deemed to have had 100% Region Availability. Any downtime occurring prior to a successful Service Credit claim cannot be used for future claims. Annual Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon EC2 SLA Exclusion (defined below).
- "Region Unavailable" and "Region Unavailability" means that more than one Availability Zone in which you are running an instance, within the same Region, is "Unavailable" to you.
- "Unavailable" means that all of your running instances have no external connectivity during a five minute period and you are unable to launch replacement instances.
- The "Eligible Credit Period" is a single month, and refers to the monthly billing cycle in which the most recent Region Unavailable event included in the SLA claim occurred.
- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon EC2 account.

Service Commitments and Service Credits

If the Annual Uptime Percentage for a customer drops below 99.95% for the Service Year, that customer is eligible to receive a Service Credit equal to 10% of their bill (excluding one-time payments made for Reserved Instances) for the Eligible Credit Period. To file a claim, a customer does not have to have wait 365 days from the day they started using the service or 365 days from their last successful claim. A customer can file a claim any time their Annual Uptime Percentage over the trailing 365 days drops below 99.95%.

Operational Availability Calculation [Misra]

- **Uptime** elements: Standby time, operating time
- **Downtime** elements
  - *Logistic*: Spares availability, spares location, transportation of spares
  - *Preventive maintenance*: Inspection, servicing
  - *Administrative delay*
    - Finding personnel, reviewing manuals, complying with supply procedures, locating tools, setting up test equipment
  - *Corrective maintenance*
    - Preparation time, fault location diagnosis, getting parts, correcting faults, testing

Example: Item-Level Sparing Analysis [Misra]

- Sparing analysis challenges
  - How many spares do you need to keep the system available at the desired rate ?
  - When are you going to need to spares (manufacturing time) ?
  - Where the spares should be kept ?
  - What system level you want to spare at ?

MTTR Examples

- Hardware MTTR with spares onsite
  - Operator available - 30min
  - Operator on call - 2 hours
  - Operator available during working hours - 14h
  - Without spares - at least 24h
- SW MTTR with watchdog
  - Reboot from ROM - 30s
  - Reboot from disk - 3 min
  - Reboot from network - 10 min

MTTR << MTTF [Fox]

- Armando Fox on ‚Recovery-Oriented Computing‘
  - A = MTTF / (MTTF + MTTR)
    - 10x decrease of MTTR as good as 10x increase of MTTF ?
    - MTTF‘s are not claimable, but MTTR claims are verifiable
    - Proving MTTF numbers demands system-years of observation and experience
    - Lowering MTTR directly improves user experience of one specific outage, since MTTF is typically longer than one user session
  - HCI factor of failed system
    - Miller, 1968: >1sec “sluggish”, >10sec “distracted” (user moves away)
    - 2001 Web user study: ~5sec „acceptable”, ~10sec „excessively slow“

MTTR << MTTF [Fox]

- Proposal: Utility curve for recovery time
  - Factors: Length of recovery time, level of service availability during error state
  - Key distinction between interactive (session-based) and non-interactive systems
- If error state leads to some steady-state latency
  - For how long will users tolerate temporary degradation ?
  - How much degradation is acceptable ?
  - Do they show a preference for increased latency vs. worse QOS vs. being turned away and incentivized to return?
- Long recovery times are often reasoned by stateful components
  - Utilize alternative architecture concepts

Availability

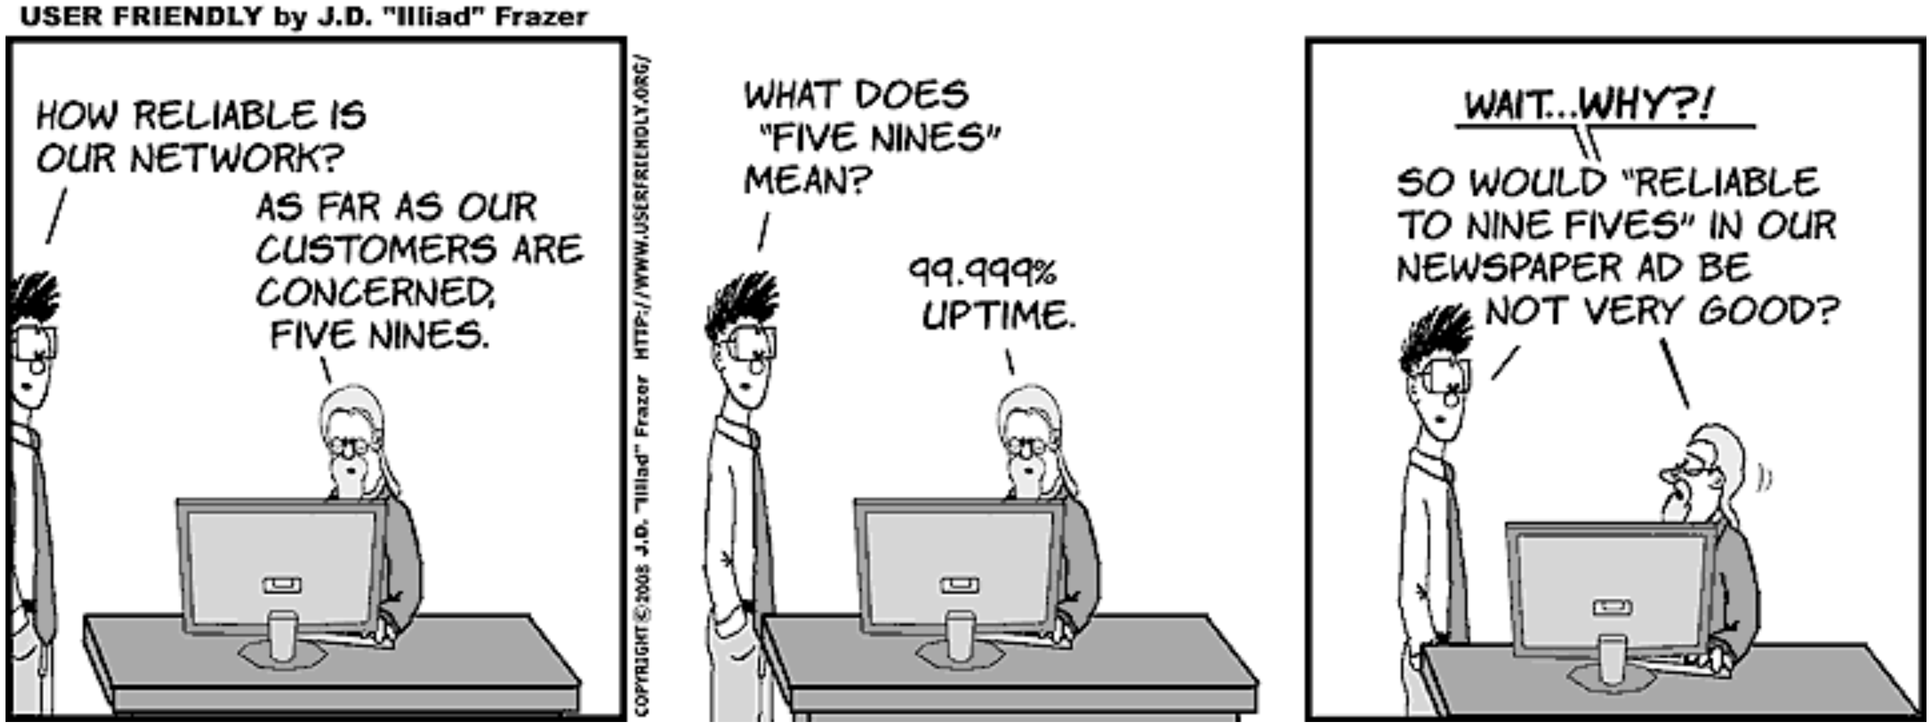
USER FRIENDLY by J.D. "Illiad" Frazer
HOW RELIABLE IS OUR NETWORK?
AS FAR AS OUR CUSTOMERS ARE CONCERNED, FIVE NINES.
COPYRIGHT©2008 J.D. "Illiad" Frazer HTTP://WWW.USERFRIENDLY.ORG/
WHAT DOES "FIVE NINES" MEAN?
99.999% UPTIME.
SO WOULD "RELIABLE TO NINE FIVES" IN OUR NEWSPAPER AD BE NOT VERY GOOD?
WAIT...WHY?!

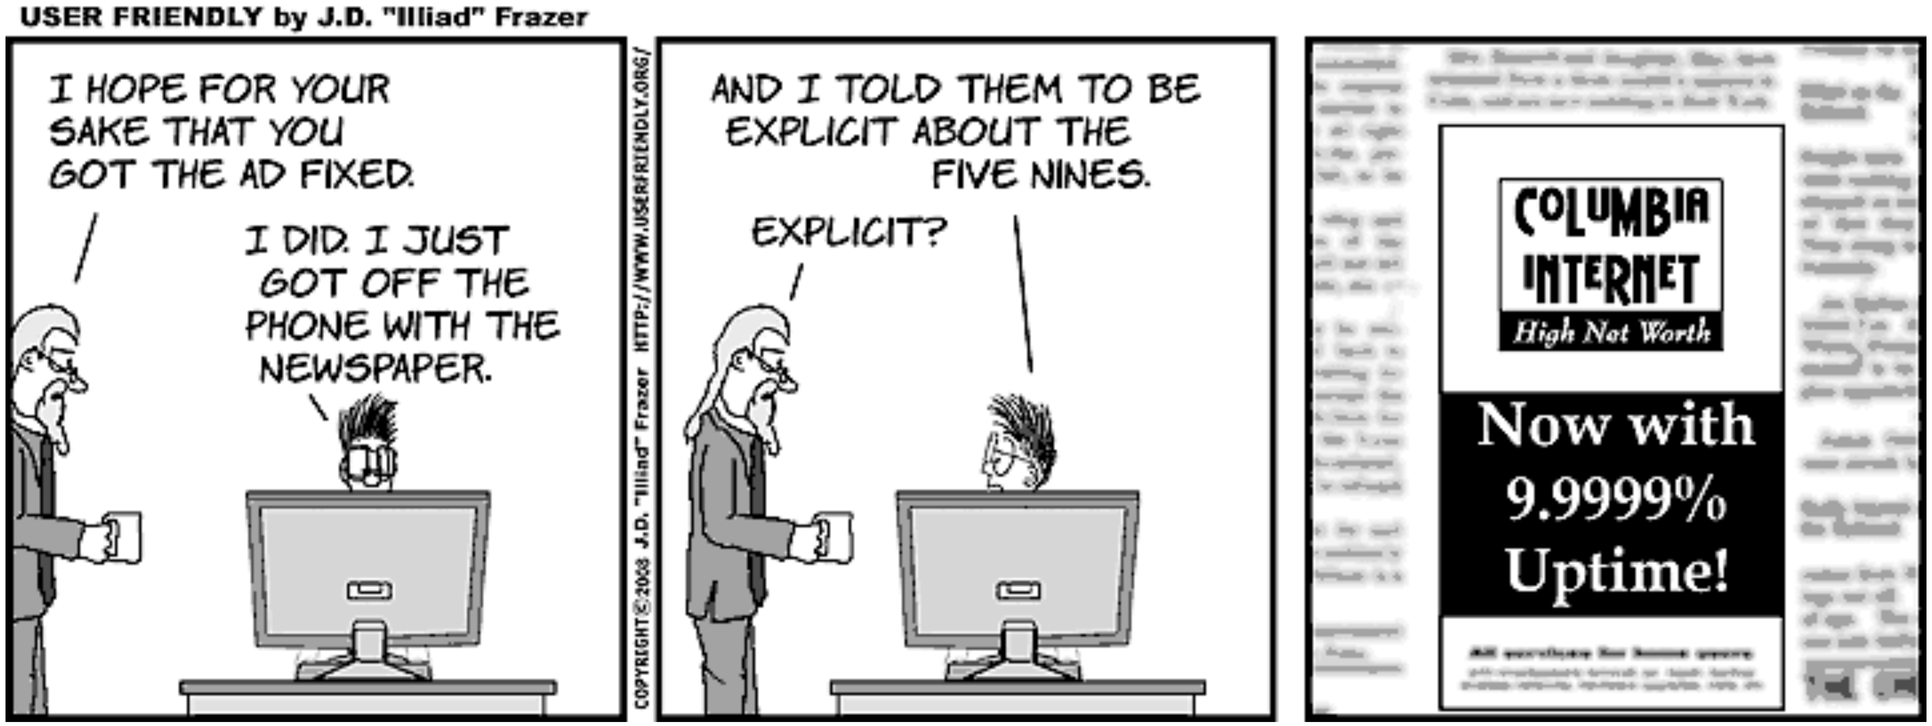
USER FRIENDLY by J.D. "Illiad" Frazer
I HOPE FOR YOUR SAKE THAT YOU GOT THE AD FIXED.
I DID. I JUST GOT OFF THE PHONE WITH THE NEWSPAPER.
COPYRIGHT©2008 J.D. "Illiad" Frazer HTTP://WWW.USERFRIENDLY.ORG/
AND I TOLD THEM TO BE EXPLICIT ABOUT THE FIVE NINES.
EXPLICIT?
COLUMBIA INTERNET
High Net Worth
Now with 9.9999% Uptime!