Reliability centric design for optimum availability of network elements

Elmer Tolentino
Foundation Technologies, Juniper Networks

Abstract: The technological evolution of packet based network elements (NEs) has required the necessity of meeting reliability objectives set by legacy circuit switched NEs. In order to achieve these objectives, the NE design must focus and understand early, reliability centric design principles. This paper will define current reliability expectations set by customers and with practical judgment, optimized reliability tasks to meet those expectations with design cost considered. System reliability engineering will discuss predictions, system modeling, testing, and field metrics.

Keywords: High availability (HA), mean time between failure (MTBF), system reliability analysis (SRA), reliability block diagram (RBD), Markov model, accelerated life test (ALT), infant mortality stress screening, and monthly return rate (MRR).

1. Key terms defined

Key terms used in this paper are first defined. *Reliability* most often is the probability of an item to perform its intended task without failure for a desired period of time in specified environments and given confidence. A more comprehensive definition for reliability can be found in [1]. *Availability* is the degree to which an item is in an up-state when the unit is called for operation at any random time. A *Network Element* (NE) is a stand-alone system such as a router, switch, hub, security device, or any other device related to processing, routing, filtering, or forwarding network traffic. The NE *system* comprises of sub-assemblies such as, to name a few, a power supply, processor card, and backplane.

2. The cost of downtime

System design of a NE for high reliability requires first an understanding of the requirements of the target market. Legacy circuit switched NEs, like a traditional Class 5 voice switch, have been required from service providers to have the capability to support 5-nines (99.999%) high availability (HA). Its unavailability ties directly with the criticality of downtime.
If, for example, a service provider network handled Emergency 911 calls, a single dropped call due to a network outage could be life threatening. Consider as well a network handling a large exchange company like NASDAQ which electronically processes up to 16,000 transactions per second [2]. Downtime of a single minute can lose a substantial amount of fees. In consequence to the service provider, NASDAQ will recoup costs based on contractual penalty obligations.

3. Legacy NEs reliability demands

Packet based NEs have risen with asymptotic behavior driven by the Internet revolution. Legacy circuit switched NEs have established requirements for high reliability which newer technologies must meet. Service providers care to a lesser degree on what technology is deployed for a given service. They are concerned more with the total cost of ownership which includes support costs after initial purchase costs. The demands for reliability must then meet, or moreover, exceed what legacy NEs have already achieved.

Reliability centric system design principles will be discussed that when implemented meet such customer demands. System reliability engineering activities will discuss the reliability cost relationship, setting reliability objectives, predictions and analysis, system modeling, reliability testing, and field metrics.

4. Reliability cost relationship

Any NE can be designed for high reliability if costs could be ignored. This behavior, however, can limit the end profit potential and be detrimental to sales. Figs. 1 and 2 show the desired regions of reliability from a supplier's view point and from a customer's view point, respectively. For suppliers, designing a new NE should consider that too high reliability will increase the end selling price, which may reduce market demand. Too high a selling price, when compared with lower priced competition, may lose sales even if the product is reliably superior. In contrast, if reliability is too low, the supplier's after shipment costs (warranty, repair depot, goodwill, etc.) rise, again increasing the supplier's selling price. Customers, in consequence, total costs rise with low reliability due to an increase in support costs (support personnel, maintenance, spares reserves, downtime, etc.).

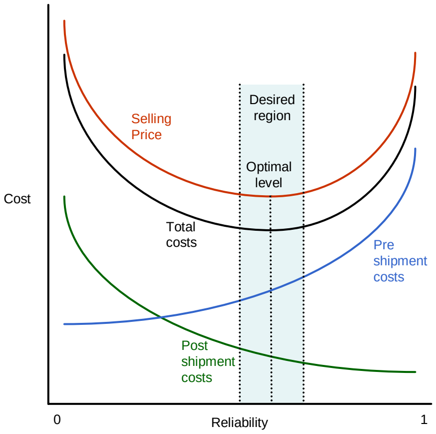

Fig. 1. Reliability cost relationship – Supplier's viewpoint

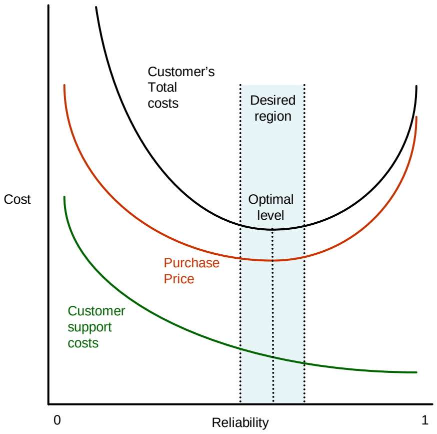

Fig. 2. Reliability cost relationship – Customer's viewpoint

The exercise in determining the costs set aside specifically for reliability, requires insight in NE applications. NEs residing in the core of a network are more likely to require higher budgets since downtime can impact thousands of users at a time. NEs residing at the edge may require less since its failure effects are typically local, affecting only the users connected to its network interfaces. Commercial NEs residing at the home or small office may require even less. However, it is duly noted that if it is anticipated that volume shipments will be high, aiming for higher levels of reliability must be considered. A 10% annual failure rate of an NE with 10,000 install base, yields 1,000 field returns for that year.

Another case lies with the importance of high revenue generating NEs. Regardless of where the NE may reside in a network, if it will be a major role of a company's revenue stream, higher levels of reliability must be considered. It will guarantee customer satisfaction, which can translate to higher sales, and solidify customer loyalty. In this special case, increasing a reliability budget to the higher end is fundamentally acceptable.

5. Reliability objectives and customer expectations

First and foremost, a NE must perform its intended task. System design must not be limited by the demands for high reliability. It is the role of the reliability engineer to work with the designer so that reliability expectations can be met. Outside the scope of this paper but also important is the software design and when coupled with hardware, guarantees that the software is of high quality, defect-free.

Reliability objectives should be set based on customer expectations of reliability parameters. The most common reliability parameters requested are MTBF and availability. As previously mentioned, network-core NEs are expected to be HA. The current desired level of availability has risen to 6-nines (99.9999%) or less than 0.5 minutes per year downtime. The interface blades that connect to the NE have less stringent levels and are typically 4-nines (99.99%) or less than 50 minutes per year downtime per interface. 4-nines availability level is also expected of 1RU to 3RU type NEs (e.g. 24-port Ethernet switch) where the severity of failure only impacts a small set of users.

From past experience with large service providers, customer MTBF expectations are high irrespective of technological complexity. Customers will always compare the MTBF of a previous design and expect as good or better MTBF. It is then best to analyze market expectations, previous technology's

MTBFs as published by suppliers and aim to meet or exceed these MTBFs.

MTBF is still a widely debated topic and is often misunderstood with customers and suppliers. However, the conversion to failure rate is not. A customer for a field replaceable unit (FRU) of high complexity (thousands of individual components) may, for example, expect 0.5% monthly return rate (MRR). This translates into 146 khrs MTBF using Eq. 1. Other failure rate conversions can be found in [8], Appendix A.

$$MTBF = \frac{730}{MRR} \qquad \textbf{\textit{Eq. 1}}$$

Note that for less complex FRUs, the complexity of the hardware design is usually simple. This results in lower component count, less power, and less heat dissipation, leading to higher MTBF. Simplification is often over-looked and understated. If a complex design can be segmented into functional FRUs, each performing a specific task, not only will the MTBF be higher per FRU but the maintenance is also optimized. A switch fabric, for example, separated into 4+1 redundant FRUs also enjoys the ability to gracefully degrade during failures. Only until the second fabric failure will total NE throughput capacity be degraded by 1/4th the total bandwidth. Segmentation also works well with port density; e.g., using two FRUs, 20 DS3 ports each, to support 40 DS3 ports. On an individual basis, each FRU now has double the MTBF. Please see Section 6 for further discussion on MTBF.

The system reliability design to achieve 6-nines availability requires all common functions to be redundant or have extremely high reliability. A core router is an example of a NE requiring 6-nines. The common functions that can cause an outage include power, backplane, switch fabric, routing engine, and control card. Failure of any of these functions would cause the entire NE to be down and hence, should be fully redundant.

The level of redundancy requires an in-depth study of the hardware requirements. Control and route functions typically are 1:1 redundant. Power and switch fabric can be 1+1 or N+1 redundant. A backplane can be designed with redundant paths so that no single path can cause an outage. Although, if the backplane were to fail, how could it be repaired without powering down? Due to this backplane repair outage scenario, it should be further designed to be as passive as possible. There should be no active components requiring repair.

The maintainability aspect of an NE is rather simple. In short, all pluggable cards should be hot-pluggable and any initialization or re-start event should be as short as possible, if not removed completely. The ability to hot-plug is of high importance to availability since it reduces the number of outages. For example, if a redundant control card failed and it was *not* hot-pluggable, the NE would have to be powered down to swap it out which is a total NE outage.

The time to re-start is function dependent. If the back-up routing engine, for example, takes 5 minutes to become fully operational after primary failure, this may not meet the 6-nines objective. In that case, the software design must establish protocols to speed up the switchover process. If a software upgrade requires a reboot, that is another outage event that can impact availability. The reliability plan should consider software development costs related to developing in-service software upgrade functionality to address this case.

Pizza box and stand-alone type NEs are a special instance where availability is less stringent. A customer understands that the level of internal redundancies will be low and maintainability also low. Typically, only the power supply would be a FRU. Nonetheless, even these NEs can be designed for NE-level redundancy, if the market has expectations for it. Stackable NEs, for example, are 6-nines capable due to the communication design between peers which allow for the synchronization of common state information, a failover protocol design, and mastership intelligence.

6. Achieving reliability goals

Early in system design, functional block diagrams are already available that graphically depict the concepts required to perform an intended task. Ideas of hardware component requirements should also be available. At this early stage, it is an important principle to analyze design reliability. Will the product feasibly meet MTBF and availability expectations within reasonable cost? Low MTBF equates to low reliability, causing the total cost for both suppliers and customers to rise. Availability if not achieved results in a product that can not even be deployed.

To determine the MTBF of a new design with no historical field data, reliability prediction models are often used. For NEs, large service providers most often request the use of Telcordia's SR-332, *Reliability Prediction Procedure of Electronic Equipment*. This standard has become one of the most widely used standards in the networking industry mainly because the derivative of the base failure rates is from networking and telecommunication applications.

Each individual component is first analyzed to determine the black box steady-state failure rate, λ_{BB}, as shown in Eq. 2.

$$\lambda_{BB} = \lambda_G \pi_Q \pi_S \pi_T \tag*{\textbf{\textit{Eq. 2}}}$$

λ_G = generic steady-state failure rate
π_Q = quality factor
π_S = electrical stress factor
π_T = temperature stress factor

After each component steady-state failure rate is predicted, the “Parts Count” method in which the assembly failure rate is assumed to be equal to the sum of each individual component is used as shown in Eq. 3.

$$\lambda_{PC} = \pi_E \sum_{i=1}^{n} N_i \lambda_{SS_i} \tag*{\textbf{\textit{Eq. 3}}}$$

λ_{PC} = parts count steady-state failure rate
n = number of different components in the assembly
N_i = quantity of i^{th} component type
π_E = unit environmental factor

An important relationship to understand is the effect of temperature stress and component failure rate. In 1889, Sweedish physicist and chemist Svante August Arrhenius empirically demonstrated chemical reaction rates, using sucrose inversion that increased exponentially with temperature. Electronic components have been observed to behave in this same manner with respect to failure rate and temperature. The relationship, $\lambda(T)$, is shown in Eq 4.

$$\lambda(T) = e^{\frac{-qE}{KT}} \tag*{\textbf{\textit{Eq. 4}}}$$

q = electron charge (1.60206 x 10^{-19} C)
E = activation energy (eV)
K = Boltzman’s constant (1.38044 X 10^{-23} J/K)
T = Kelvin temperature

When comparing the failure rate of an assembly at differing temperatures (assuming E = 0.5), it shows that it approximately doubles, 2x acceleration, for every 10°C rise in temperature. Since E is failure mechanism dependent, Fig. 3 has been shown. For values of E, please see SR-332 section 9.1.

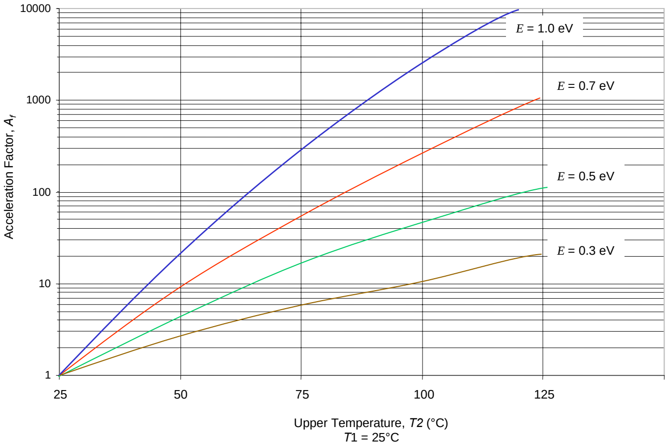

Fig. 3. *Acceleration factor and temperature of varying E*

When an assembly does not meet MTBF goals, the components causing the increased failure rate can be further analyzed. The initial analysis should review the base failure rate as derived from Telcordia. Research can be done based on each component manufacturer for HTOL (high temperature operating life) data which provide the component FIT rate. Normally, the FIT results are performed at the upper 60% confidence limit (CL) based on the Chi-square distribution. The MTBF and Chi-square relationship is shown in Eq. 5 and Ref. [3] shows Chi-square values. Note, if the board prediction is using 90% CL, the HTOL FIT rate should be re-calculated at 90% CL as well. For RAM, ROM, CPU, and optical transceivers most often HTOL data is available.

$$MTBF = \frac{2t}{\chi^2_{1-\alpha}}\text{, evaluated at DF} \tag*{\textbf{\textit{Eq. 5}}}$$

$MTBF = 1/FIT$
t = Total cumulative test time
χ^2 = Chi square value at $1-\alpha$ and DF
α = confidence limit
DF = $2r$ = degrees of freedom
r = (number of failures) +1

Another avenue for MTBF improvement can be to use industrial grade parts. They have shown to exhibit higher MTBF especially in the case of hard drives. Industrial greade HD MTBF has been found on average to be 1.2 Mhrs while commercial grade is 500 khrs. Improved cooling performance can reduce internal operating temperature. As shown earlier, a 10°C reduction in temperature can double MTBF. For power components, derating electrical stress below rated values can further reduce failure rate.

If this secondary analysis is completed and MTBF goals are still not met, upper management should decide if the project should continue referring back to costs. If the project does proceed, a risk mitigation plan should be in place to properly support reliability. Such a plan should consider increased spare reserves, field failure rate trending, additional costs for HA software features development, required failure analysis for field failures, additional reliability testing, legal contractual changes, etc. Once field performance proves field failure rate to be acceptable, these mitigation steps can be removed or lessened to a degree as deemed comfortable by upper management.

To determine availability of a new design, reliability modeling using reliability block diagrams (RBDs) and Markov state-space diagrams should be used that define the operational behavior of the NE. It is again suggested that standards based methods be used such as Telcordia's SR-1171, *Methods and Procedures for System Reliability Analysis (SRA)*. As defined in the standard, the steps in system reliability modeling and analysis are:

1. System definition
2. Measures selection
3. Architecture analysis and decomposition
4. Architecture modeling
5. Model parameter determination
6. Model solution
7. Model parameter sensitivity analysis
8. Analysis summary

System definition requires a complete understanding of how each assembly within a NE work together to perform an intended task. It is at this step where RBDs are defined and redundancy schemes established. Fig. 4 shows an RBD example for a NE.

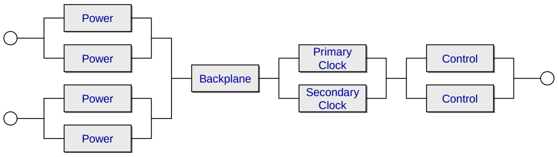

Fig. 4. *Reliability block diagram of an NE*

Measures selection simply defines what will be predicted. For NEs, customers typically want to know availability and downtime. Architecture analysis and decomposition is the analysis of individual sub-assemblies within the NE with respect to the measures that are to be predicted. The complexity of decomposition can be eased by performing a failure modes effects analysis (FMEA) [4]. Architecture modeling involves creating reliability models which define the system. Markov modeling is typically done here. Fig. 5 shows a Markov model example for a cooling system in a fixed NE.

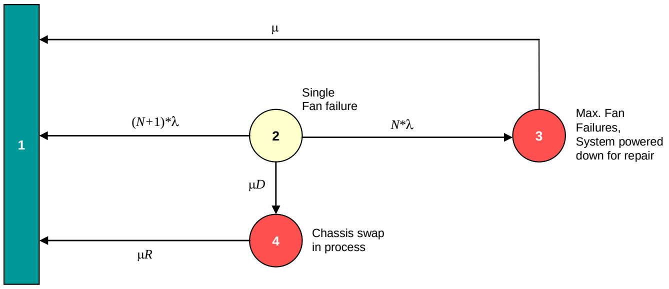

Fig. 5. *Cooling system Markov model*

Model parameter determination defines failure rates, repair rates, detection probabilities, etc., used in modeling. From Fig. 5, N, λ, and μ are examples of parameters that need to be defined. Solving the models should be aided by software tools to prevent from errors and ease complexity. When solving a Markov model, numerical methods are used with a novel approach described by Trivedi [5].

Model parameter sensitivity analysis is used when availability goals are not met. What parameters within the models can be changed so that availability goals are achieved? Is there a particular NE assembly that if designed for redundancy can result in a 6-nines NE? Will reducing the failover time from 5-minutes to 10 seconds for redundant processors allow us to meet the availability objective? It is at this step in an SRA where the power of design change through sensitivity analysis can allow a NE to meet availability objectives.

Lastly, when the design is nearing completion, MTBF and availability estimates should be analyzed again. In this phase, final BOMs (bill of materials) are used, system architecture is finalized, software tests have initial estimates for failover, etc. Also, Telcordia can then be used to review all results and formerly publish a reliability analysis report verifying the findings. This is an advantageous step in reliability analysis since it uses a 3rd party to review and ensure all reliability estimates are in compliance with Telcordia standards. From experience, when a large service provider recognizes Telcordia publication of a NE reliability analysis, the results are fully accepted.

7. Reliability testing

The discussion on reliability testing is hardware centric focusing on the functional and environmental requirements as expected by customers. In order for a product to be reliable, it must first perform its intended task within the specified environments. DVT (design verification test), signal integrity, and environmental tests are typically done here.

DVT is usually done using the 4-corner's test which varies voltage and temperature as shown in Fig. 6. This test must be accompanied by supporting diagnostics so that failures during test can be recognized immediately. Signal Integrity tests are typically classified into two sections, Signal quality and Timing budget. Signal quality tests check for signal amplitude, ringing, undershoot and overshoot. Timing budget tests check for potential timing violations by measuring the signals set-up & hold and comparing it with the component's specifications.

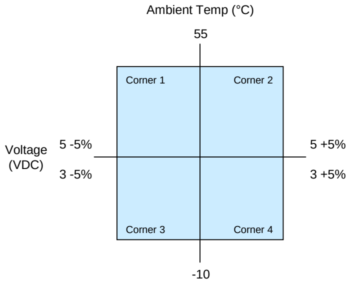

Fig. 6. *DVT 4-Corner test cases*

Operational temperature and humidity test follow Telcordia's GR-63-CORE, *NEBS Requirements: Physical Protection*. Again, large service providers prefer standards based methods and GR-63 has been widely accepted. Based on the standard, NEs are expected to be subjected to varying degrees of temperature, temperature change rates, and humidity as shown in Fig. 7. Customers expect the NE to operate reliably in these conditions throughout the NE's steady-state life.

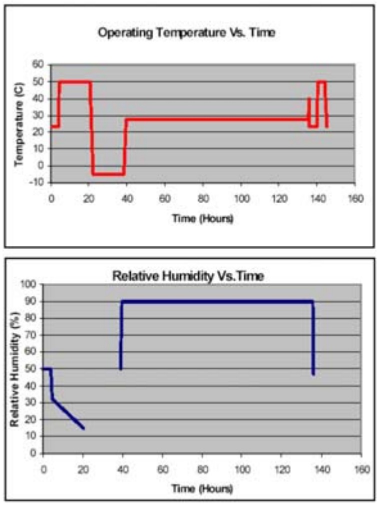

Fig. 7. *Operating temperature and humidity test profile*

To confirm operation at varying altitudes, the GR-63 Altitude test is performed. This test verifies operation of up to 13,000 feet above sea level. The shipping environment will also subject the NE to excessive stresses that can cause failure. GR-63 provides test methods for transportation and storage environments, equipment handling, packaged shock, and unpackaged shock.

Mechanical stresses encountered during earthquakes and office vibration must also be tested to verify operation in such conditions. Telcordia has performed a study on earthquake stresses and have found that for each risk zone (see Table 1. and Fig. 8), there is a 90% likelihood that an earthquake event of this severity will not be exceeded over a 50-year period. Hence, survivability under these stresses increases the confidence that the NE will operate reliably under such conditions throughout the NE's steady-state life.

Earthquake Risk Zone	Richter Magnitude	Low Frequency Ground Acceleration (Gs)	Low Frequency Upper Building Floor Acceleration (Gs)
1	< 4.3	< 0.05	< 0.2
2	4.3 – 5.7	0.05 – 0.1	0.2 – 0.3
3	5.7 – 63	0.1 – 0.2	0.3 – 0.4
4	6.3 – 7.0	0.2 – 0.4	0.4 – 0.6
5	7.0 – 8.3	0.4 – 0.8	0.6 – 1.0

Table. 1. *Earthquake risks and zones*

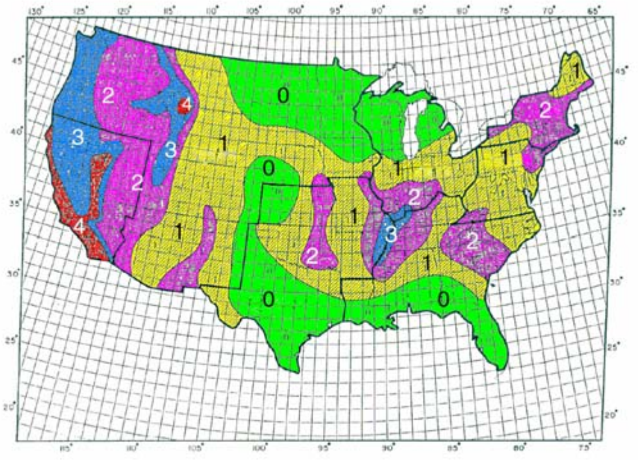

Fig. 8. *Earthquake zone map*

Hygroscopic dust testing verifies that sulfate and nitrate salts found in air do not cause failures. This test case simulates the electrical effects of fine mode particle dust by reducing surface insulation resistance. Humidity levels are then increased and the NE must remain functional after contamination

An accelerated life test (ALT) allows for the simulation of field operating hours in a shortened test time by using stresses much higher than found during normal conditions. The first ALT test to be discussed is Gaseous Contaminants. This test subjects the NE to the 95th percentile; meaning 95% of the time, the level of contaminant is lower. The test is run for 10 days which equates to 15 years of operating life.

HALT (highly accelerated life test) rather stimulates, instead of simulate, an environment through extreme stress testing to find failure modes not found in previous tests. HALT theory assumes that increasing a given stress is equal to allowing that product to operate at its normal stress for a longer period of time. Note, however, HALT is not a reliability demonstration test. MTBF can not be deduced from results found in HALT. No studies to date have shown that failure modes found in HALT will be the exact primary failure mode that will be found in the field. Nevertheless, customers do request HALT and the determination of when a NE will fail is very good information.

For example, if a product's operating specification was 40°C and through other environmental testing the maximum temperature tested was 50°C, there is no idea of when the NE will break. Is it 51°C or 80°C? A product with an 11°C operating margin is obviously less reliable with respect to temperature than the NE that has a 40°C margin. HALT testing takes the NE to the fundamental limits of technology (FLT). If the NE does not break prior to that, the HALT test is deemed successful. If the NE does break but then the hardware is fixed so that the FLT is reached, that too is a successful HALT test. The key test elements in HALT are:

1. Cold thermal step stress
2. Hot thermal step stress
3. Rapid thermal cycling
4. Quasi-random vibration
5. Combined environments

A combined environment example profile is shown in Fig. 9.

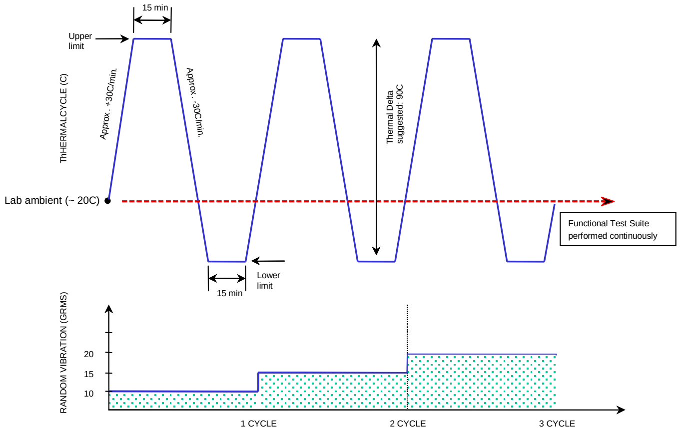

Fig. 9. *Combined environments HALT test profile*

Similar to DVT, diagnostics play an integral role in determining if the NE has failed. Also, it is extremely important that thermally sensitive components be monitored by thermocouples if the NE does not support temperature monitoring. A microprocessor, for instance, at 80°C test temperature may report a 100°C case temperature which can lead to a 110°C junction temperature as depicted in Fig 10.

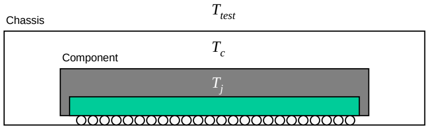

Fig. 10. *Temperature differences during reliability testing*

T_{test} = Chamber set test temperature
T_c = Chassis temperature
T_j = Component junction temperature

RDT (reliability demonstration testing) is an ALT which provides insight to actual MTBF. Based on the Arrhenius model, the test temperature is raised to reduce test time as well as increasing the number of test samples. The test costs for implementing this is high. Thermal chambers are required to support the sample size. Diagnostic test equipment is required to monitor all test samples. An engineer is required to monitor daily test results for at least 1 full month and provide support for any RCA (root cause analysis). Most HALT proponents oppose RDT due to this. On the other hand, most RDT proponents oppose HALT since it may not reflect the actual failure modes found in the field. It is the view of the author that if costs permit, RDT testing should be limited to 1-2 months of testing and an MTBF goal of 75-100 khrs at 60% CL. For less complex NEs where sample size can be high, higher

MTBF goals are acceptable as long as test time is controlled. If upper management deems it necessary, 90% CL can be tested upon completion.

The test must be started on, or before, first customer shipments. There is little benefit to start RDT for a product that already has years of field install base with enough data to estimate field MTBF. Some reliability programs implement ORT (on-going reliability test) which continuously validates MTBF from a rotating test sample (e.g. 1 week-in, 1 week-out, new units 1 week-in, etc.). This is done, for the most part, to monitor the manufacturing plant since most suppliers outsource the manufacturing process. The cost for implementing ORT is also high and is often not feasible.

Infant mortality failures are typically quality and manufacturing related. NEs will usually undergo a 24-72 hour burn-in at 40-50°C. When statistical evidence proves that failures are found much shorter (e.g. 97% of failures are found before 8 hours of burn-in) than the burn-in test time can be optimized. At more cost, HASS (highly accelerated stress screen) can reduce test time. Most often this is not a feasible stress screen since mechanical fixtures need to be designed to support vibration testing. Thermal cycling has found to be effective in removing weak solder joints but its cost is also higher than traditional burn-in.

Software reliability testing is likewise as important but is outside the scope of this paper. It is worthy to note however that steps in modeling software reliability is still emerging but that industry experts have not yet come to a common standards based methodology for the assessment of software reliability. Most customers expect software to be defect free upon product receipt, and the effect to overall failure rate is zero. Of course, that is not true and in the field, defects (bugs) in the software will occur that may lead to an outage event depending on its severity and other factors.

An interested customer typically allows a supplier to field trial the NE in their network in a controlled test environment. This is also known as a Beta test. This collaboration with the customer is a very effective tool in reducing defects when a NE finally goes live in an active network. It is difficult for a test lab to simulate the customer's exact network environment. Hence, Beta test opportunities should be highly valued.

8. Field metrics

The traditional bathtub curve plots failure rate over time and shows different regions in a NEs product life cycle as shown in Fig. 11.

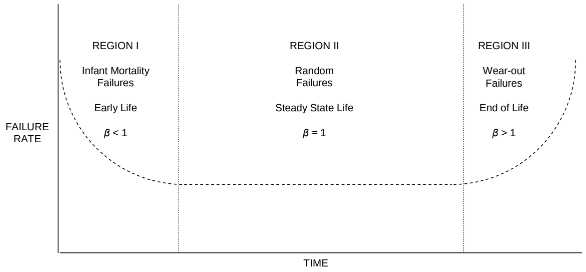

Fig. 11. *Reliability bathtub curve*

Region II is steady-state and is the area where MTBF can be estimated. Plotting MRR over time is useful in assessing field reliability performance and for plotting the bathtub curve. Region I and III can be analyzed using Weibull when shape parameter, β, are <1 and >1, respectively [6].

Customer contracts often dictate monthly return rate thresholds that if exceeded, require immediate attention and may include penalties. Suppliers will track the monthly return rate to abide by this requirement. Also, suppliers with high reliability in mind will also set field reliability goals based on predictions. If the field reliability falls below the predicted MTBF, the supplier will pro-actively initiate RCA. Other flagging mechanisms can also be set in place for cases like increasing failure rate trends, return rate spikes, etc. Fig. 12 shows an example of MRR over time. In this example, the long-term average MRR is approximately 0.5% or 146 khrs MTBF. The spike in Month 8 would trigger RCA and Months 9-12 verifies that the corrective action has resolved all issues, restoring the steady-state failure rate trend.

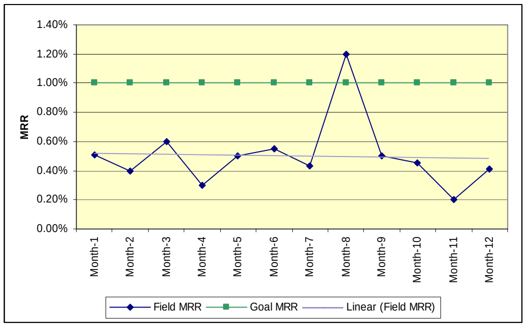

Fig. 12. *Example FRU field MRR chart*

This very simple concept of tracking monthly return rate addresses customer contracts and consistently monitors field reliability performance. The sustaining engineering team typically will lead the efforts in field reliability support and improvement. A closed-loop failure reporting and corrective action system (FRACAS) is a key element for sustaining [7].

9. Conclusion

Customer requirements force new technologies to reach for higher levels of reliability as established by legacy NEs. New NE systems designed for high reliability must set targets, do due diligence for reliability design analysis, allow for change when reliability objectives are not met, manage risk, plan to test, and follow through with field metrics. From past experience with NEs, when these reliability practices are consistently used, the probability that field reliability will exceed predictions is greater than 90%.

10. References

[1] D. Kececioglu, *Reliability Engineering Handbook*, Volume 1, Prentice-Hall, pg. 5, 1991.
[2] Hewlett-Packard Development, *NASDAQ posts big gains with HP information lifecycle management solution*, HP press release, Jan. 2005.
[3] R. Dovich, *Reliability Statistics*, ASQC Quality Press, pg. 80, 1990.
[4] J. Bowles, *Fundamentals of Failure Modes and Effects Analysis*, Tutorial Notes of the annual reliability and maintainability symposium (RAMS), Jan. 2008.
[5] K. Trivedi, *Probability & Statistics with Reliability, Queuing and Computer Science Applications*, Prentice-Hall, eqs. 8.13, 8.14, and 4.39, 1982.
[6] R. Abernathy, *The New Weibull Handbook*, 2nd Edition, Jun. 1996.
[7] P. O'Conner, *Practical Reliability Engineering*, 4th Edition, Wiley and Sons, pgs. 333-334, 2002.
[8] SR-332, *Reliability Prediction Procedure for Electronic Equipment*, Issue 2, Telcordia Technologies, Sept. 2006.
[9] SR-TSY-001171, *Methods and Procedures for System Reliability Analysis*, Issue 1, Telcordia Technologies, Jan. 1989.
[10] GR-63-CORE, *NEBS Requirements: Physical Protection*, Issue 3, Telcordia Technologies, Mar. 06.
[11] GR-2840-CORE, *Generic Requirements for Environmental Stressing Applied to Telecommunications Products*, Issue 1, Telcordia Technologies, Jun. 1995.
[12] W. Ireson, C. Coombs, R. Moss, *Handbook of Reliability Engineering and Management*, 2nd Edition, McGraw-Hill, pgs. 12.16-12.20.
[13] S. Ali, *Digital Switching Systems System Reliability and Analysis*, McGraw-Hill, 1998.
[14] B. Dodson, D. Nolan, *Reliability Engineering Bible, The Complete Guide to the CRE*, Quality Publishing, 1995.
[15] G. Hobbs, *Accelerated Reliability Engineering, HALT and HASS*, Wiley and Sons, 2000.
[16] G. Hobbs, *Advanced HALT and HASS, Highly Accelerated Life Tests used in Design, Highly Accelerated Stress Screens used in Productions*, Tutorial notes, Dec. 1994.
[17] G. Bastien, *Advanced ESS Technical Seminar*, Tutorial notes, Jul. 1996.
[18] F. Jensen, N. Petersen, *Burn-in, An Engineering Approach to the Design and Analysis of Burn-in Procedures*, Wiley and Sons, Aug. 1990.
[19] E. Amerasekera, F. Najm, *Failure Mechanisms in Semiconductor Devices*, 2nd Edition, Wiley and Sons, 1997.

Elmer Tolentino has over 13 years of reliability engineering experience specializing in technologically advanced networking systems. He was a design team member of the very first Juniper Network's router and of Marconi's (now Ericsson) very first ATM switch. He has worked on over 70 different networking systems and over 250 different types of sub-system components, which include network interfaces (up to OC-768), route processors, switch fabrics, power and cooling systems, fiber optic transceivers, clocking systems, and other supporting subsystems. He has authored over 20 system reliability & availability analysis reports published by Telcordia Technologies. He received the B.S. degree in mechanical engineering from San Jose State University in 1995 and is pursuing a Ph.D. in reliability engineering. He is currently employed by Juniper Networks as a Staff Reliability Engineer, leading system reliability design and analysis for all Juniper hardware products, and is a senior member of ASQ (CRE, 2001) and member of IEEE.