The Distribution of Time to Recovery of Enterprise IT Services

Ulrik Franke, Hannes Holm, and Johan König

Abstract—**The context of this article is the availability of enterprise IT services, a key concern for many enterprises. While there is a plethora of literature concerned with service availability, there is no previous systematic empirical study on IT service time to recovery following outages. The existing literature typically assumes a distribution, or builds on analogies to related areas such as software engineering. Therefore, our objective is to find the statistical distribution of IT service time to recovery. Method-wise, this investigation is based on logs of more than 1 800 incidents in a large Nordic bank, corresponding to more than 11 000 hours of recorded downtime. Five possible distributions of time to recovery from the literature were investigated using the Akaike Information Criterion to find the distribution offering the best fit. The results show that the log-normal distribution outperformed the others for all tested service channels (collections of IT services). It is concluded that the log-normal distribution offers the best fit of IT service time to recovery. Using this distribution in simulation and decision-support tools offers the prospect of better predictions of downtime and downtime costs to the practitioner community.**

Index Terms—**Enterprise IT services, incident logs, log-normal distribution.**

Abbreviations & Acronyms

AIC	Akaike Information Criterion
CDF	Conditional Distribution Function
CI	Confidence Interval
COCOMO	Constructive Cost Model
EXP	Exponential Distribution
GAM	Gamma Distribution
IT	Information Technology
LN	Log-normal Distribution
LTLN	Laplace Transform of the Log-normal Distribution
PDF	Probability Density Function
SLA	Service Level Agreement
TOGAF	The Open Group Architecture Framework
WBL	Weibull Distribution

Notation

AIC_i	The AIC value of a model i being evaluated
$AIC_{\min}$	The AIC value of the model with the lowest AIC
Δ_i	$AIC_i - AIC_{\min}$
K	The number of parameters in a model
n	The size of a sample

I. Introduction

With increasing market competition and increasing reliance upon information systems, Information Technologies (IT) is becoming ever more important to business operations [14]. Unfortunately, this importance also entails an increased sensitivity to failing IT services and IT systems downtime. While the *Encyclopedia of Information Assurance* contains numerous entries on Business Continuity Planning, it also acknowledges that “Business resumption and disaster recovery planning is probably the part of information security that is easiest to overlook and postpone” [25]. Nevertheless, continuity planning is becoming increasingly important following the growth of the Internet and e-business [26], and is stressed in practitioner frameworks such as the Information Technology Infrastructure Library (ITIL) framework [54]. Renowned consultancy Gartner regularly produces reports not only on how to assess IT service availability levels [51], but also on how to calculate the costs of downtime [35], and how to assess the cost of maintaining continuous service availability [36]. Furthermore, these issues receive a lot of management attention. For example, in a 2010 survey of Chief Executive Officers’ and business executives’ top 10 IT uncertainties, the reliability of IT comes second [24]. It is also instructive to consider media reporting on IT outage incidents. When Bank of America suffered an online banking outage in January 2011, this event brought about a landslide of bad publicity, including reminders of previous unavailability incidents [10]. The outage itself lasted for approximately ten hours, but the consequences lasted much longer.

In today’s business environment, where IT is increasingly procured as a service regulated by Service Level Agreements (SLAs), detailed knowledge of statistical distributions for outage durations is becoming more valuable. Recent technology trends such as cloud computing further underscore the importance of availability risk management [42], but the need to properly balance the cost for achieving high availability against the cost of unplanned outages remains a key challenge also for companies that maintain their own data centers [6].

Manuscript received March 03, 2013; revised October 11, 2013 and January 30, 2014; accepted March 19, 2014. Date of publication August 06, 2014; date of current version November 25, 2014. Associate Editor: S. Shieh.

U. Franke and H. Holm are with FOI—the Swedish Defence Research Agency, SE-164 90 Stockholm, Sweden (e-mail: ulrik.franke@foi.se; hannes.holm@foi.se).

J. König is with the Department of Industrial Information and Control Systems, KTH—the Royal Institute of Technology, SE-100 44 Stockholm, Sweden (e-mail: johan.konig@ics.kth.se).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TR.2014.2336051

In the insurance business, actuarial data are a crucial asset. Insurance companies can go out of business if they do not use correct probability distributions for the hazards they insure. As the IT service provision market matures, service providers who have agreed to pay fines for lengthy service outages will be prudent to obtain similar statistical knowledge. How service-providers should best allocate their resources to maximize revenue while not violating availability SLAs is an active topic for academic research [7], [23]. Nevertheless, companies still have poorly characterized availability objectives, and IT departments struggle to express availability so that it makes sense to the business side [45]. Recent research suggests that IT decision-makers assigning availability SLAs might be less than rational in their risk management [18], and that the information currently used in SLAs can lead to suboptimal decisions [31]. However, as the field of availability management becomes more mature, perhaps turning into a systematically developed "service level engineering" [30] discipline, we can expect explicit, standardized decision-making criteria for availability management to evolve. This evolution is similar to the way measures like standard deviation and value at risk (VaR) have become standard in financial reporting. This paper offers an applied availability risk management example in Section VII. The overall aim of this paper is to further this maturation, including the development of decision-making criteria that facilitate rational decision-making.

A. Scope of the Paper

This paper contributes to quantitative IT service availability management by investigating the statistical distribution of recovery times in enterprise IT services. The investigation is based on logs of more than 1 800 incidents in a large Nordic bank. Another, more exploratory contribution is an investigation into the root causes of these downtime incidents. The incidents recorded all concern IT service continuity relevant for business processes, which is precisely what enterprises such as Bank of America care about.

The scope of the article is the time to recovery of enterprise IT services. This scope is different from software repair as studied in software reliability. Whereas software repair is about finding and removing bugs in code, enterprise IT service recovery potentially concerns everything that causes service disruptions, including hardware, software, configuration, and human error. Software repair can play an important role here, but the service is then typically recovered through a roll-back to the last working version for the duration of the actual debugging effort. Once the bug has been found and removed, the software is deployed again. Hence, while software repair can take days or weeks, service recovery in critical systems often takes minutes. This distinction is further elaborated in Section IV-A.

B. Outline

The remainder of the paper is structured as follows. Section II explains the importance of knowing the proper distribution of time to recovery. Section III contrasts the present contribution with some related work, followed by some method considerations in Section IV. Section V introduces the data set. Section VI is the locus of the main contribution. Here, the results of the studies are described, followed by an applied example in Section VII, illustrating the potential financial impact of modeling time to recovery using the wrong distribution. A discussion on validity and reliability then ensues in Section VIII, followed by a discussion of the strengths and weaknesses of the contribution in Section IX, and some concluding remarks in Section X.

II. The Importance of the Time to Recovery Distribution

The importance of knowing the time to recovery distribution has been stressed repeatedly in the literature. For example, Snow *et al.* have simulated the chances of SLA violations, and emphasize that the tail of the repair distribution is crucial [53]. In later work, Snow and Weckman extend the argument, and warn against reasoning about availability based on averages rather than full distributions [52]. Building on this strand in the literature, Section VII goes on to practically demonstrate how the distribution of time to recovery has a financial impact for a service provider.

Marques *et al.*, who investigate how to best design SLAs, explain the importance of knowing both parameters (e.g., means) and distributions of time to recovery. While means suffice for some agreements, agreements such as having 95 percent of requests served within a certain time, require knowledge about the full distribution [37]. Similar conclusions are reached by González and Helvik [22].

To summarize, with erroneous assumptions about distributions, models will give the wrong results, and predictions will not be valid.

III. Related Work

The scope of this paper is the time to recovery of enterprise IT services. As explained above, this differs importantly from the field of software reliability. This distinction is further elaborated in Section IV-A. Nevertheless, much of our related work comes from software reliability. The following exposition shows how our investigation has been inspired by methods and results from this area, even though no previous work precisely shares our scope.

Software reliability has mostly been concerned with time to *failure* (producing famous failure rate models such as the Jelinski-Moranda model [27], the Schick-Wolverton model [49], and the Goel-Okumoto imperfect debugging model [20]). However, papers that have empirically explored the *repair* rates of various types of failures are more closely related to our work.

For example, Gokhale and Mullen [21] analyzed software defect repair times for nine different Cisco Systems product families and a total of more than 10 000 samples. Each sample consists of a time from identifying a bug in a software to creating a patch for that bug. The authors compared the data fit for the exponential distribution (EXP), the log-normal (LN) distribution, and the Laplace Transform of the log-normal distribution (LTLN). The fit of the distributions to the data was evaluated using the Akaike Information Criterion (AIC). The best fit was found for the LTLN and LN, with a slight favor towards the LTLN. However, while software defect repair rates are valuable for e.g., vulnerability discovery models such as

Alhazmi's [2], their application is less straightforward for enterprise IT services, which is our scope. That is, the translation from time to repair for software defects to time to recovery for higher-level business services is not trivial, as a software defect does not need to imply service availability issues, and vice versa.

Khoshgoftaar and Woodcock [29] applied AIC to select the best among different software reliability models for predicting the number of remaining errors, and the time to their discovery during software development and testing. This approach differs from our scope in that we consider IT services in operation, not software being developed, and that we study time to recovery, not number of remaining errors.

Schroeder and Gibson [50] studied the mean time to repair for availability issues of 20 different systems, mostly large clusters of SMP (Symmetric-Multi-Processing) and NUMA (Non-Uniform-Memory-Access) nodes at Los Alamos National Laboratory. 23 000 failures from a period of nine years were analyzed. The authors analyzed distribution fit for Weibull (WBL), gamma (GAM), EXP, and LN distributions. LN was found to provide the best fit out of the four distributions considered. While the topic certainly is relevant to the enterprise level, there are several issues that could create problems if one were to apply the results to enterprise IT services directly. (i) The formal statistical analyses are not fully characterized in the article; a quantitative comparison of the candidate distributions is missing. This quantitative comparison is the role played by AIC in our study. (ii) The data were collected over a period of nine years. It is therefore likely that several major software and hardware changes have been carried out during the time. (iii) Only 20 systems were studied, possibly for validity reasons, as this was the population over which the researchers had sufficient control to conduct a study. Nevertheless, most enterprises have hundreds, if not thousands, of systems, as reflected in our data set.

Plank and Elwasif [41] used the results of three workstation monitoring projects to study the applicability of theoretical equations concerning the performance of checkpointing. As the mathematical model assumes failure and repair to be Poisson processes, with correspondingly exponentially distributed times, part of the paper investigates this assumption. The authors show that the EXP distribution provides very low degrees of fit. No other distributions were analyzed. Similarly to Plank and Elwasif, Long *et al.* [34] carried out a survey regarding Internet host reliability, and evaluated the suitability of the Poisson process for time to repair. The authors concluded that the EXP distribution was a poor fit for repair times. Such studies, showing the poor performance of the EXP distribution, are an important motivating factor for our work, which tests multiple distributions for the best fit.

Labovitz *et al.* [33] studied network failures in two scenarios: (i) 3 years of network failures in five of the U.S. Internet routing exchange points, and (ii) 12 months of network failures in a large regional service backbone. While the paper provides cumulative distributions for the time to recovery of these scenarios, it does not explore the fit of any theoretical distributions. Also, the results are probably limited to the domain of the study, network failures on the Internet.

To conclude, while there are many authors who have explored repair rates for different kinds of failures, none of these studies is fully applicable in the enterprise IT service context.

IV. Log Data Analysis Method

A. Data on Enterprise IT Service Recovery Times

Our scope of investigation is enterprise IT services. An IT service, as defined in the ITIL framework, "is based on the use of Information Technology and supports the Customer's Business Processes. An IT Service is made up from a combination of people, Processes, and technology and should be defined in a Service Level Agreement" [55]. An enterprise, as defined in The Open Group Architecture Framework (TOGAF) standard, is the "highest level (typically) of description of an organization and typically covers all missions and functions. An enterprise will often span multiple organizations" [56]. Enterprise IT services, therefore, are Information Technology used to support business processes on the highest level of an organization.

To better understand the relation between a piece of software and an enterprise IT service, it is useful to consider the description of Johnson *et al.* [28]):

> "Although system users might sometimes feel that the systems fail to deliver, the information systems in a company are there to provide value to the business. Even when successful, however, the information systems themselves need support to continue delivering services to their users. As briefly mentioned in the previous chapter there is thus a causal flow from the IT organization through the information systems to the business [...]"

Thus, an IT service is not only about technology, but about technology in an organizational setting. Software might be a necessary precondition for a successful enterprise IT service, but it is not a sufficient one.

For the purpose of this article, we define time to recovery as the time from loss of service until it is restored to operation. This definition is similar to the definition given by Milanovic [38], and includes the times to report, locate, investigate, and repair a system, as well as the time it takes for a repair team to physically show up, and any time required to restart the system afterward. The end-to-end term is sometimes used to signal the explicit inclusion of each of these times. Following the importance of the business process in the ITIL definition of an IT service, loss of service occurs only when someone in the business takes notice. Availability of an IT system by itself is of no consequence until it reaches the business; and conversely, unavailability that does not impact a business process is not (yet) a problem.

This scope is very different from e.g., Gokhale and Mullen [21] who investigate repair times in a much more restricted software engineering sense, considering only the process of finding and repairing an actual bug. Such a debugging effort by a programmer can take a day, a week, or a month. Service recovery, on the other hand, is typically a matter of a one-minute roll-back to the last working version (and when it is not, media coverage such as that given by Charette [9] often sets in). On top of that, a software bug or a hardware fault do not necessarily entail service downtime.

To properly investigate the time to recovery of enterprise IT services, a few criteria have to be met: (i) the data must not reflect just a single technical solution or architecture, because enterprise IT services are diverse in this respect; (ii) the data must not measure a mere part of end-to-end recovery time, and (iii) the data must not reflect technical solutions rather than services (e.g., hard drive crashes rather than storage outages). Our dataset, made available from a large Nordic bank, satisfies these criteria, and will be further described in Section V.

B. Statistical Analysis

Based on the previous work presented in Section III, we consider five distributions in terms of degree of fit to recovery times of enterprise IT services: EXP, LN, LTLN, WBL, and GAM. LN, LTLN, WBL, and GAM have displayed the highest degree of fit in previous studies [21], [50], [57], [41]. It should be noted that even though these studies do not concern enterprise IT services, they are the most natural starting point available. While EXP to our knowledge has not shown good fit in any published study, it is a simple, closed-form distribution for time to recovery of enterprise applications known from the literature [17], [40], and thus imperative to research.

EXP The exponential distribution is widely used to model repair time distributions [21]. However, its popularity is mostly due to its mathematical simplicity rather than its precision [43], [44], [17].

LN A random variable X is log-normally distributed if its natural logarithm is normally distributed, i.e., $\ln X \in \mathbf{N}(\mu, \sigma)$. The LN distribution is often used as a more realistic model of repair times [43].

LTLN The Laplace transform of the LN distribution is based on work by Mullen, including Gokhale and Mullen [21]. It assumes that the *time to repair* is determined by a randomly drawn *repair rate*, λ, which itself follows an LN distribution. It is thus a doubly stochastic model. With the repair time being exponentially distributed conditioned on a given repair rate λ, the probability that the defect of rate λ is not repaired until time t or later is $e^{-\lambda t}$. It follows that the probability that the defect is repaired before time t is $1 - e^{-\lambda t}$. Knowing that λ is LN distributed, the cumulative distribution function (CDF) of the time to repair can be obtained by the following integral.

$$\begin{aligned} M(t) &= 1 - \int_{\lambda=0}^{\infty} e^{-\lambda t} dL(\lambda) \\ &= 1 - \int_{\lambda=0}^{\infty} e^{-\lambda t} \frac{1}{\lambda\sigma\sqrt{2\pi}} \quad e^{-\frac{(\ln\lambda-\mu)^2}{2\sigma^2}} d\lambda \end{aligned} \tag{1}$$

This integral is equivalent to the Laplace transform of the LN CDF, and has no simple form. A numerical method for solving it, used in this paper, is given in [39].

TABLE I
EMPIRICAL SUPPORT FOR MODELS USING AIC DIFFERENCES; RULES OF THUMB REPRINTED FROM [5]

Δ_i	Level of empirical support of model i
$0 - 2$	Substantial
$4 - 7$	Considerably less
> 10	Essentially none

WBL The Weibull distribution is known for its ability to describe a number of phenomena, and is sometimes used to model the repair times of general systems [13], [8], [58]. By adjusting two parameters a and c (both being positive real numbers), the WBL distribution can be parameterized into various other probability distributions.

GAM The gamma distribution is a continuous function with two adjustable parameters p and a (both being positive real numbers), that is sometimes used to model time to repair [50], [57].

1) The Akaike Information Criterion: To compare the relative statistical goodness of fit of the distributions to the data, the AIC, a standard technique to rank alternative models, was used. The AIC was introduced to extend the method of maximum likelihood estimation to the situation of multimodel choice [1]. As pointed out by Burnham and Anderson [5], Akaike's criterion links Boltzmann's entropy, Kullback-Leibler information, and maximum likelihood, thus tying together information theory with statistics. Essentially, the AIC is an estimator of the expected relative Kullback-Leibler information [5].

Conceptually, the AIC can be seen as adding a penalty to models with many parameters, thus rewarding not only fit but also simplicity [21]. The AIC is a measure of the badness of fit of a model (larger AIC means worse model) defined with parameters estimated by the maximum likelihood method [1]. It is defined as

$$AIC = -2 \cdot \text{log likelihood} + 2 \cdot \text{number of parameters}. \tag{2}$$

The lower its AIC, the better a distribution fits the data. To evaluate distributions against each other, the difference $\Delta_i = AIC_i - AIC_{\min}$ is formed, where AIC_i is the value of the model being evaluated, and $AIC_{\min}$ is the value of the model with the lowest AIC. By definition, $AIC_{\min}$ thus has $\Delta_i = 0$. Some rough rules of thumb given by Burnham and Anderson [5] are given in Table I.

There are many other rules of thumb, all expressing essentially the same verdict. For example, Sakamoto *et al.* [47] maintain that, if the difference of AIC between models is larger than $1 \sim 2$, then the difference is considered to be significant. According to Gokhale and Mullen [21], a difference of four is considered very significant.

If there are too many parameters (K) in relation to the size of the sample (n), AIC may perform poorly [5]. According to Burnham and Anderson [5], the ratio n/K should be above 40 for AIC to perform well. In our case, we have removed a few small sub-datasets to consistently meet the $n/K > 40$-rule, as described in the next section. The K used was 2, because this value is the largest number of parameters used in any of the statistical distributions considered, thus giving a sample size threshold of 80.

V. Data

The analysis is based upon incident data from a major Nordic bank. The full dataset consists of 2 335 incidents recorded from January 2009 to May 2011. In all, these incidents correspond to 709 558 minutes, or almost 12 000 hours of recorded downtime.

The incidents are categorized by the bank in a number of ways:

- by *service* affected, i.e., which one out of several hundreds of distinct IT services that was down;
- by *department* affected, i.e., which IT department was responsible for the IT service that was down;
- by *cause*, i.e., an unstructured textual description of the cause of the downtime;
- by *impact*, i.e., a classification into the categories Down, Partly down, No impact, or Planned, along with some blank (i.e., not classified) entries; and
- by *channel*, i.e., a business side categorization of services.

The analysis in Section VI builds upon the categorization of services into *channels* such as Automated Teller Machines (ATMs), Internet banking, credit card payments, etc. A channel contains a number of inter-related IT services working in concert to offer these more coarse-grained business services. There are 14 different channels available when categorizing an incident. These channels are pre-defined by the business side, not by the IT department. Even though some services participate in several channels, incidents are reported channel-wise. Whenever a service incident is reported for channel A, only channel A is affected. Other channels can still use the service, and if they cannot then it is reported as a separate incident. All categorizations are done by the team (typically 2 or 3 people) working to resolve the incident. If they are uncertain about how to categorize something, they go back and check how similar incidents have been categorized historically. In this sense, the categorization is reliable over time. The only category where the members of the response team are free to describe the problem on their own without pre-defined templates or categories is the description of the cause of the downtime.

The categorization by impact is important because the scope of the paper is to study the time to recovery of enterprise services. Therefore, incidents recorded as in the Planned or No Impact categories were removed from the data set, as were incidents with downtime recorded as having durations of 0 minutes (these 0 minute incidents largely coincided with the Planned and No Impact incidents).

However, incidents where the service was recorded as Partly Down were included. The reasons are twofold. (i) On a strict reading, partly down means that something is wrong, and that there is indeed an outage of some kind. (ii) More importantly, however, the incidents are recorded with a start time, a stop time, a problem description, and a resolution following recovery efforts, thus fulfilling reasonable requirements to constitute times to recovery in the sense relevant to be measured, and therefore are included in the analysis.

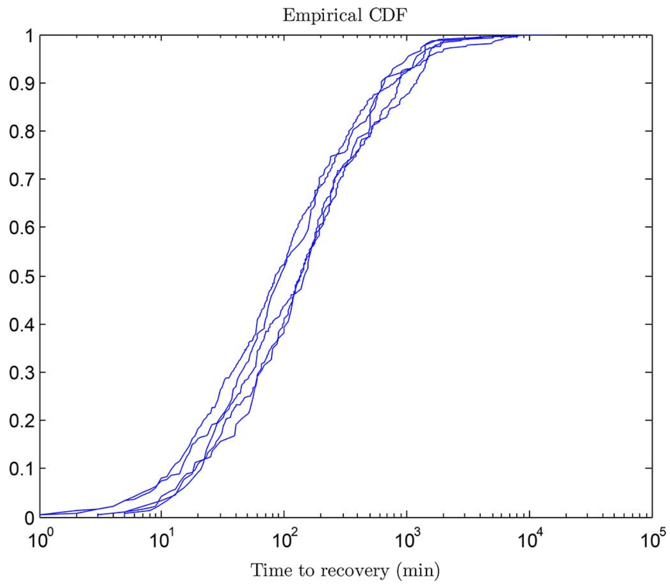

Fig. 1. CDFs for the recovery times of the 5 channels (x-axis in minutes).

TABLE II
Percentage of Down-Time in 2010 for Two Channels, Distributed Over 14 Root Causes

	Channel 2		Channel 4	
	Minutes	Percentage	Minutes	Percentage
Physical environment & Infrastructure redundancy	5	0.01%	912	1.00%
Requirements and procurement	3295	5.43%	0	0.00%
Operations	10422	17.16%	32021	35.07%
Change control	20277	33.39%	19983	21.89%
Technical solution of backup	0	0.00%	1966	2.15%
Process solution of backup	205	0.34%	0	0.00%
Data & Storage architecture redundancy	190	0.31%	1933	2.12%
Internal application failures	15929	26.23%	12159	13.32%
External services that fail	10140	16.70%	8285	9.07%
Network redundancy	18	0.03%	0	0.00%
Network failures	5875	9.67%	405	0.44%
Physical location	0	0.00%	0	0.00%
Resilient client/server solutions	2952	4.86%	9294	10.18%
Monitoring of the relevant components	661	1.09%	9455	10.36%
Other	682	1.12%	15019	16.45%

Furthermore, some of the channels contained too few incidents to pass the $n/K > 40$-rule of [5]. Out of the 14 channels, 9 were removed from further consideration as each of them contained fewer than 80 incidents over the 2.5 year span. The remaining 5 channels together contains 1 876 incidents, corresponding to 672 272 minutes, or just over 11 000 hours of recorded downtime. The channels analyzed are themselves composed of 322 individually named services, provided by particular IT systems. Fig. 1 depicts the cumulative distribution functions for the recovery times of the 5 channels investigated.

Unfortunately, as described above, the causes are not categorized into distinct, well-defined categories in the bank reporting system. Therefore, to gain an understanding of the causes, a new categorization of incidents had to be made. 14 different categories from Franke *et al.* [19] (where full definitions are also given) were used for taxonomy, and a categorization of a subset of the data was made with the assistance of the bank. The subset thus selected for manual categorization consisted of all incidents that occurred from January to December 2010 in channels 2 and 4. This data set corresponds to 364 incidents, totaling 152 031

TABLE III
MAXIMUM LIKELIHOOD PARAMETER ESTIMATES FOR DIFFERENT CHANNELS AND DISTRIBUTIONS

Channel	λ (EXP)	μ (LN)	σ (LN)	μ (LTLN)	σ (LTLN)	a (WBL)	c (WBL)	p (GAM)	a (GAM)
1	0.003152	4.51918	1.568718	-5.04408	1.130726	202.0314	0.633572	0.510346	621.7372
2	0.002531	4.799661	1.638889	-5.35148	1.134507	271.73	0.650241	0.533271	740.8047
3	0.003416	4.689199	1.360495	-5.17643	0.935257	217.9972	0.707761	0.620953	471.4182
4	0.002305	4.929999	1.492937	-5.44585	1.045473	293.8988	0.659519	0.548046	791.7077
5	0.002702	4.937629	1.370357	-5.44216	0.863051	275.8705	0.714759	0.628647	588.7442

TABLE IV
DISTRIBUTION COMPARISON BY AIC VALUE FOR CHANNEL RECOVERY TIMES. THE LN DISTRIBUTION SHOWS THE BEST FIT IN ALL CASES

Channel	EXP	LN	LTLN	WBL	GAM	n
1	11588.38	10952.68	10984.05	11103.19	11244.93	857
2	5320.007	5118.015	5120.244	5150.254	5190.779	381
3	2019.132	1940.628	1953.605	1975.136	1993.439	151
4	5631.945	5375.747	5394.32	5451.042	5510.229	398
5	1232.658	1190.552	1194.942	1207.479	1219.261	89

TABLE V
EMPIRICAL SUPPORT FOR DIFFERENT DISTRIBUTIONS. $\Delta_i < 2$ GIVES SUBSTANTIAL SUPPORT, $4 < \Delta_i < 7$ CONSIDERABLY LESS, AND $\Delta_i > 10$ ESSENTIALLY NONE, ACCORDING TO [5]

Channel	Δ_{EXP}	Δ_{LN}	Δ_{LTLN}	Δ_{WBL}	Δ_{GAM}
1	635.7068	0	31.37301	150.5072	292.2508
2	201.9924	0	2.229666	32.23886	72.76411
3	78.50434	0	12.97756	34.50862	52.81074
4	256.1982	0	18.57285	75.29468	134.4823
5	42.10582	0	4.39031	16.92771	28.7094

TABLE VI
95% CONFIDENCE INTERVALS FOR THE PARAMETER ESTIMATES μ AND σ OF THE LN DISTRIBUTION

Channel	95% CI for μ (LN)	95% CI for σ (LN)	n
1	[4.414004, 4.624356]	[1.497804, 1.646733]	857
2	[4.634571, 4.964751]	[1.530197, 1.76433]	381
3	[4.470435, 4.907962]	[1.222417, 1.534017]	151
4	[4.782878, 5.077119]	[1.395925, 1.604549]	398
5	[4.64896, 5.226298]	[1.194382, 1.607629]	89

minutes, or approximately 2 500 hours of recorded downtime, i.e., roughly a fifth of the entire material. The distribution of root causes per channel is given in Table II. It should be noted that the percentages do not sum to unity, because the categories are not fully mutually exclusive. (Downtime can have several causes. First, a system may go down because of an internal application failure, and then fail to come back up due to inadequate monitoring. Such downtime can reasonably be attributed to both factors.)

VI. RESULTS

The maximum likelihood estimates of the parameters of the EXP, LN, LTLN, WBL, and GAM distributions for each of the 5 channels are given in Table III.

The distributions are compared in terms of AIC in Tables IV and V. The LN distribution shows the best fit, i.e., the lowest AIC, in all cases.

Using the rules of thumb for empirical support given by Burnham and Anderson [5], we can conclude that no alternative to the LN distribution gets any substantial support. The EXP, WBL, and GAM distributions all get essentially no support in any of the five channels. The LTLN distribution also gets essentially no support in three of the five channels, whereas in one channel it gets considerably less support and in another it gets somewhere between substantial and considerably less. It is clear that all the channels in the dataset considered are best described by the LN distribution.

Even though the focus of this article is to find a statistical distribution suitable for modeling enterprise IT service recovery times, it is interesting also to say something about the parameter values of the estimated distribution.

Table VI gives the 95% confidence intervals for the parameter estimates μ and σ of the LN distribution from Table III. As is to be expected, the intervals are quite narrow for the larger data sets.

VII. AN APPLIED EXAMPLE

This section gives a small example that illustrates how the distribution of time to recovery has a financial impact for a service provider.

Consider a service provider about to sign an SLA. One of the provisions in the SLA stipulates that a fine must be paid whenever the time to recovery of a service exceeds h hours. How should the service provider determine a reasonably reliable h?

First, of course, the service provider needs to set its risk level (accounting for the size of the fine, among other things). Assume that the acceptable risk level is deemed to be that a fraction r (for risk) of outages are handled within the stipulated time frame (e.g., 95%). If the distribution of time to recovery is known, finding h then corresponds to solving the following integral equation

$$r = \int_0^h f_X(x)dx. \tag{3}$$

As before, f_X is the probability density function, which depends upon distribution specific parameters. In some simple cases, such as that of EXP, h can be found explicitly:

$$r = \int_0^h \lambda e^{-\lambda x} dx = 1 - e^{-\lambda h} \Rightarrow h = -\frac{\ln(1-r)}{\lambda}. \tag{4}$$

Other distributions, such as LN, require numerical methods.

TABLE VII
DIFFERENT h VALUES (IN HOURS) FOR CHANNEL 1, SHOWING THE IMPACT OF MODELING USING DIFFERENT DISTRIBUTIONS AT DIFFERENT RISK LEVELS r. RECALL THAT THE PARAMETERS USED ARE ALL MAXIMUM LIKELIHOOD ESTIMATES MADE FROM THE SAME DATA SET

	Channel 1 [hours]		
$r =$	90%	95%	99%
EXP	12.2	15.8	24.4
LN	11.4	20.2	58.8
LTLN	11.5	19.3	50.0
WBL	12.6	19.0	37.5
GAM	14.2	20.2	34.7

Solving (3) using the parameters of Table III for Channel 1 (the channel with the most samples) yields h, for different r levels, as illustrated in Table VII.

Table VII exhibits some interesting features. First of all, the number of hours h required to reach the desired risk level r increases sharply as r grows for all distributions, as expected. But, second, this increase depends a lot on the distribution of time to recovery. For example, LT and LTLN are virtually indistinguishable at the 90% level, but differ markedly at the 99% level. Third, the rank order of the distributions also changes as r grows. While GAM requires the most hours to meet the 90% risk level, it is surpassed by all but EXP at the 99% level. EXP itself has a much lighter tail than the others, and requires less than half the number of hours to reach the 99% risk level than does LN. It is clear that a service provider who signs an SLA using the EXP assumption, but whose time to recovery is actually LN distributed, will need to pay far more fines than expected. These findings are consistent with the importance of recovery time distribution tails stressed by the work cited in Section II.

The importance of correct distribution assumptions is further emphasized by the decreasing margins of service providers. Renowned consultancy Gartner recommends service providers to continuously analyze their competitive position, and protect margins [46]. In some sectors, such as communications, the pressure on profit margins even calls for rethinking whole business paradigms to remain competitive [15]. In this environment, every dollar of profit margin is important, and service providers need solid quantitative decision support, including reliable distributions of time to recovery, to build robust business cases.

VIII. VALIDITY AND RELIABILITY OF FINDINGS

Overall, *validity* refers to "the best available approximation to the truth of falsity of propositions" [11]. While there are many facets of validity that can be discussed, the arguably most important to address concern internal validity, external validity, construct validity, and reliability [3], [4]. The validity of the findings from the present study are discussed along three of these dimensions in the three next subsections. *Internal validity*, the truth value that can be assigned to the conclusion that a cause-effect relationship between a statistically independent variable and a statistically dependent variable has been established within the context of a particular research setting [4], is left out on purpose as the study does not evaluate any causal relationships. It is not within the scope of the paper to investigate how different variables affect recovery times, as we study the goodness-of-fit for different statistical distribution models.

A. External Validity

External validity refers to the generalizability of causal findings with respect to the desired population and settings [4].

An argument for the validity of the study is the fact that it is based upon the very metrics that are of business relevance to the enterprise. The logs examined concern precisely the service downtime noticed by customers and staff in their work, i.e., concrete cost-driving downtime and service recovery times. If these downtimes were not valid, such as if they did not measure something that is interesting and relevant to measure, the bank would surely discontinue their measurement.

A more complex issue has to do with selection bias in the population of enterprises, and the final verdict on the validity of this study might have to wait until the analysis has been repeated on different sets of data in future studies.

B. Construct Validity

Construct validity refers to the extent to which an identified causal relationship can be generalized from the particular methods and operations of a specific study to the theoretical constructs and processes they were meant to represent (rather than the desired population and settings, as for external validity) [4].

In terms of operationalized theories, this study covers two areas: statistical goodness-of-fit, and time to recovery. While the prior did not require any adaption (merely use of standard statistical techniques), the operationalization of time to recovery could be discussed. In the study, downtimes were recorded as either Down, Partly Down, Planned, or No Impact. This study operationalizes time to recovery as incidents causing either complete (Down) or partial (Partly Down) unavailability, and exclude incidents that were planned (Planned) or yielded no impact (No Impact) on the availability of a service. These choices are reasonable in the sense that they reflect recoveries from significant incidents (rather than scheduled maintenance or incidents without impact). However, the actual business impact of incidents corresponding to these categories (especially Partly Down) is not completely certain. Findings from the study should be considered in the light of these design choices.

C. Reliability

Reliability is the extent to which an instrument produces consistent or error-free results [3].

One potential reliability concern is the reliability of the mapping of incidents to the right channels. However, there are two reasons to believe that this categorization is reliable. First, the classification is made by 2 or 3 experts, not by the users themselves. Second, in case of doubt, the experts use previous classifications as guidance. Thus, the classification should be regarded as robust over time.

Another potential reliability concern is the reliability of the recovery times. The times are set manually, and thus subject to human error. However, as noted above, the staff managing the incident logs is small and professional, and they also use a validation system to avoid simple data entry errors. Furthermore, as the statistics are used for monthly evaluation and reporting, large aberrations from data entry errors are likely to be detected.

IX. Analysis and Discussion

As pointed out in Section IV-A, the enterprise IT service recovery scope of this paper is different from the software repair scope in much of the literature. This difference is very clearly illustrated by the different x-axis scales in Fig. 1 (minutes), and the corresponding Fig. 6 in Gokhale and Mullen [21] (measured in days).

That said, from an enterprise IT service availability perspective, software repair is still an important activity to understand, as it is often related to the recovery of IT services. For instance, a bug in the software supporting an online banking system can cause unavailability, more or less intermittent, of the supported enterprise IT services until the error has been repaired by the developers. Consequently, the results by Gokhale and Mullen [21] should be viewed as *complementary* to those described in this paper, and enterprise decision-makers should consider both to enable well-informed availability estimates and decisions.

While root cause analysis is not the primary subject of this paper, the root cause mapping shown in Table II does offer some preliminary insights. It is worth noting that the two top categories (Operations, and Change control) are also among the top causes identified previously in an expert survey [17]. However, other top causes from the expert survey such as Requirements and Procurement are not reflected as prominently in Table II. It may well be the case that more abstract causes such as Requirements and Procurement are rarely reflected in logs. Measuring such causes probably requires other research methods that reflect the earlier phases of system life-cycles, not only the operations phase.

How can availability be improved and outage risks properly managed? Knowing the distribution of time to recovery is not in itself sufficient. This knowledge needs to be incorporated into decision-making, as illustrated for example in Section VII.

The importance, and relative ease, of availability modeling for enterprise information systems is demonstrated by Närman *et al.* [40], where it is shown that accurate availability estimates (precision within eight hours downtime per annum) of relatively complex enterprise services can be achieved with no more than 20 man-hours of work. Investing in such modeling seems very worthwhile. Using the LN distribution, the precision of modeling frameworks such as Närman's [40] (which builds on EXP) can be further improved.

However, with the advent of cloud computing, not all companies can model their enterprise architecture in detail. Rather, the SLA is their only interface to their IT services, which is management by contract, as it is called by Sallé [48]. In these contexts, there is a need for careful availability modeling beforehand, to avoid costly mistakes. Such decision-making needs to account for the different requirements of different businesses, e.g., whether (i) more but shorter or (ii) fewer but longer outages ought to be preferred [16]. As a general rule for availability risk management, the business side and the IT department need to communicate, so that the business consequences of different IT decisions are clearly understood and quantified. Enterprise architecture models, useful for both communication and analysis [32], can play an important role here, as can IT service management frameworks such as ITIL [54].

X. Conclusions and Future Work

Previous research on availability in the enterprise setting has often assumed EXP or LN distributions of service recovery times. However, only rarely have such assumptions been justified with empirical data. Instead, analogies have been made to quite different systems such as grid computers or to software engineering, where bug fixing has been put on a par with service recovery.

In this paper, incidents causing downtime in five enterprise IT service channels in a major Nordic bank were investigated. In all, more than 1 800 incidents, corresponding to more than 11 000 hours of recorded downtime, were evaluated to see whether the recovery times conform to any of the distributions from the literature. Based on this investigation, recovery times of enterprise IT services seem to be best described by the LN distribution.

This work opens several avenues for further research. One is to repeat the analysis on different sets of data. Does the LN distribution describe service recovery times in other enterprises as well? Another direction has to do with the actual parameters of the LN distribution. What are the factors determining μ and σ? Do the parameters differ (significantly) between different lines of business, different technical architectures, or even different SLAs (as one should hope, but not take for granted)? The root cause investigation, reflected in Table II, offers another interesting perspective. Are μ and σ different for different root causes? Sometimes, though not always, the root cause is a bug that needs to be fixed. Here one might investigate whether there is a correlation between the time to recovery (this study) and the time to bug repair (e.g., [21]).

It is also instructing to look at incident data as a strategic asset. Properly used, incident data could not only help development teams to take corrective actions to avoid future outages, but also to manage uncertainty at the beginning of projects, and to develop enterprise risk management. Here, more research is needed both conceptually (what is possible?) and empirically (what is the current state of the practice in companies?).

Incident management can serve as a tool for learning, if done properly. For example, the ITIL framework prescribes the following incident management steps [54] (p. 60): Problem detection, problem logging, categorization, prioritization, investigation and diagnosis, create known error record, resolution, and closure. It would be interesting to investigate whether maturity in this process is reflected in a quicker time to recovery, or perhaps in a smaller standard deviation. Learning from incidents has been explored using simulation models [12]; empirical data would add another dimension.

For the business community, the results offer the prospect of better predictions of downtime, and downtime cost. By implementing mathematical models with more realistic assumptions about service recovery times, better decision support can be achieved.

ACKNOWLEDGMENT

The authors are grateful to the bank that made the data available for analysis. Clearly, this piece of research would not have been possible otherwise. Furthermore, the constructive comments of the anonymous reviewers helped us to substantially improve the paper, as did the remarks by Göran Ericsson and Nicholas Honeth, who generously allotted the time to read several drafts of the manuscript.

REFERENCES

[1] H. Akaike, "Factor analysis and AIC," *Psychometrika*, vol. 52, 1987.

[2] O. Alhazmi, Y. Malaiya, and I. Ray, "Measuring, analyzing and predicting security vulnerabilities in software systems," *Comput. Security*, vol. 26, no. 3, pp. 219–228, 2007.

[3] M.-C. Boudreau, D. Gefen, and D. W. Straub, "Validation in information systems research: A state-of-the-art assessment," *MIS Quart.*, pp. 1–16, 2001.

[4] M. B. Brewer, "Research design and issues of validity," in *Handbook of Research Methods in Social and Personality Psychology*. Cambridge, U.K.: Cambridge Univ. Press, 2000, pp. 3–16.

[5] K. Burnham and D. Anderson*, Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. New York, NY, USA: Springer-Verlag, 2002.

[6] D. J. Cappuccio, Ensure Cost Balances Out with Risk in High-Availability Data Centers, Gartner, Inc., Tech. Rep., Feb. 2013.

[7] E. Casalicchio, D. A. Menascé, and A. Aldhalaan, "Autonomic resource provisioning in cloud systems with availability goals," in *Proc. 2013 ACM Cloud and Autonomic Computing Conf.*, 2013, p. 1.

[8] C. Cassady, L. Maillart, R. Bowden, and B. Smith, "Characterization of optimal age-replacement policies," in *Proc. IEEE 1998 Annu. Reliability and Maintainability Symp.*, 1998, pp. 170–175.

[9] R. Charette, Power Outage at Barclays Bank Causes Chaos Saturday Afternoon in the UK, Oct. 2010, IEEE Spectrum "Risk factor" blog [Online]. Available: http://spectrum.ieee.org/riskfactor/computing/it/power-outage-at-barclays-bank-causes-chaos-saturday-afternoon-in-the-uk

[10] R. Charette, Bank of America Suffered Yet Another Online Banking Outage, Jan. 2011, IEEE Spectrum "Risk factor" blog [Online]. Available: http://spectrum.ieee.org/riskfactor/telecom/internet/bank-of-america-suffered-yet-another-online-banking-outage-

[11] T. D. Cook and C. S. Reichardt*, Qualitative and Quantitative Methods in Evaluation Research*. Beverly Hills, CA, USA: Sage, 1979, vol. 1.

[12] D. L. Cooke and T. R. Rohleder, "Learning from incidents: From normal accidents to high reliability," *Syst. Dynam. Rev.*, vol. 22, no. 3, pp. 213–239, 2006.

[13] A. De Almeida and G. Bohoris, "Decision theory in maintenance decision making," *J. Qual. Maint. Eng.*, vol. 1, no. 1, pp. 39–45, 1995.

[14] W. Delone and E. McLean, "The DeLone and McLean model of information systems success: A ten-year update," *J. Manage. Inf. Syst.*, vol. 19, no. 4, pp. 9–30, 2003.

[15] K.-Y. Foong, J. Forsman, J.-C. Delcroix, C. Patrick, A. K. Sharma, and W. L. Hahn, Predicts 2011: CSPs Must Rethink Business Paradigms to Meet Market Challenges, Gartner, Inc., Tech. Rep., Nov. 2010.

[16] U. Franke, "Optimal IT service availability: Shorter outages, or fewer?," *IEEE Trans. Netw. Service Manage.*, vol. 9, no. 1, pp. 22–33, Mar. 2012.

[17] U. Franke, P. Johnson, J. König, and L. M. von Würtemberg, "Availability of enterprise IT systems—An expert-based Bayesian framework," *Softw. Qual. J.*, vol. 20, no. 2, pp. 369–394, 2012, DOI: 10.1007/s11219-011-9141-z.

[18] U. Franke, M. Buschle, and M. Österlind, "An experiment in SLA decision-making," in *Economics of Grids, Clouds, Systems, and Services*. New York, NY, USA: Springer, 2013, pp. 256–267.

[19] U. Franke, P. Johnson, and J. König, "An architecture framework for enterprise IT service availability analysis," *Softw. Syst. Model.*, pp. 1–29, 2013, DOI: 10.1007/s10270-012-0307-3.

[20] A. L. Goel and K. Okumoto, "A Markovian model for reliability and other performance measures of software systems," in *Proc. Int. Workshop Managing Requirements Knowledge*, 1979, vol. 0, p. 769.

[21] S. Gokhale and R. Mullen, "A multiplicative model of software defect repair times," *Empir. Softw. Eng.*, vol. 15, pp. 296–319, 2010, DOI: 10.1007/s10664-009-9115-y.

[22] A. González and B. Helvik, "Guaranteeing service availability in SLAs; a study of the risk associated with contract period and failure process," in *Proc. IEEE Latin-American Conf. Communications, 2009 (LATINCOM'09)*, 2009, pp. 1–6.

[23] H. Goudarzi, M. Ghasemazar, and M. Pedram, "SLA-based optimization of power and migration cost in cloud computing," in *Proc. 2012 12th IEEE/ACM Int. Symp. Cluster, Cloud and Grid Computing (CCGrid)*, 2012, pp. 172–179.

[24] K. Harris, Gartner CEO and Senior Business Executive Survey, 2010: Perceptions of IT and Tactical Fixes, Gartner, Inc., Tech. Rep., Mar. 2010.

[25] K. Henry*, Business Continuity Planning: Case Study*. New York, NY, USA: Taylor & Francis, ch. 42, pp. 344–350 [Online]. Available: http://www.tandfonline.com/doi/abs/10.1081/E-EIA-120046814

[26] C. B. Jackson*, Business Continuity Planning: Evolution in Response to Major News Events*. New York, NY, USA: Taylor & Francis, ch. 46, pp. 377–383 [Online]. Available: http://www.tandfonline.com/doi/abs/10.1081/E-EIA-120046817

[27] Z. Jelinski and P. Moranda, "Software reliability research," *Statist. Comput. Perform. Eval.*, pp. 465–484, 1972.

[28] P. Johnson and M. Ekstedt, Enterprise Architecture: Models and Analyses for Information Systems Decision Making, Studentlitteratur, 2007.

[29] T. M. Khoshgoftaar and T. G. Woodcock, "Software reliability model selection," *Qual. Rel. Eng. Int.*, vol. 8, no. 5, pp. 457–469, 1992.

[30] A. Kieninger, J. Westernhagen, and G. Satzger, "The economics of service level engineering," in *Proc. 2011 44th Hawaii International Conf. System Sciences (HICSS)*, 2011, pp. 1–10.

[31] A. Kieninger, D. Straeten, S. O. Kimbrough, B. Schmitz, and G. Satzger, "Leveraging service incident analytics to determine cost-optimal service offers," in *Proc. 11th International Conf. Wirtschaftsinformatik*, 2013, pp. 1015–1029.

[32] S. Kurpjuweit and R. Winter, "Viewpoint-based meta model engineering," in *EMISA 2007: Proc. Enterprise Modelling and Information Systems Architectures*, Bonn, Germany, Oct. 2007, pp. 143–161, Gesellschaft fr Informatik.

[33] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental study of internet stability and backbone failures," in *Proc. 29th Annu. IEEE Int. Symp. Fault-Tolerant Computing, 1999, Digest of Papers*, 1999, pp. 278–285.

[34] D. Long, A. Muir, and R. Golding, "A longitudinal survey of internet host reliability," in *Proc. 14th IEEE Symp. Reliable Distributed Systems, 1995*, 1995, pp. 2–9.

[35] B. Malik, Q&A: How Much Does an Hour of Downtime Cost?, Gartner, Inc., Sep. 2009, Tech. Rep.

[36] B. Malik and D. Scott, How to Calculate the Cost of Continuously Available IT Services, Gartner, Inc., Nov. 2010, Tech. Rep.

[37] F. Marques, J. Sauvé, and J. Moura, "SLA design and service provisioning for outsourced services," *J. Netw. Syst. Manage.*, vol. 17, no. 1, pp. 73–90, 2009.

[38] N. Milanovic, Models, Methods and Tools for Availability Assessment of IT-Services and Business Processes Universitätsbibliothek, 2010, habilitationsschrift.

[39] R. Mullen, "The lognormal distribution of software failure rates: Application to software reliability growth modeling," in *Proc. 9th Int. Symp. Software Reliability Engineering, 1998*, Nov. 1998, pp. 134–142.

[40] P. Närman, U. Franke, J. König, M. Buschle, and M. Ekstedt, "Enterprise architecture availability analysis using fault trees and stakeholder interviews," *Enterprise Inf. Syst.*, 2012, DOI: 10.1080/17517575.2011.647092.

[41] J. Plank and W. Elwasif, "Experimental assessment of workstation failures and their impact on checkpointing systems," in *Proc. 28th Annu. IEEE Int. Symp. Fault-Tolerant Computing, 1998, Digest of Papers*, 1998, pp. 48–57.

[42] S. Prakash, "Risk management: Cloud computing considerations," *Canadian Manage. Account.*, no. 2, p. 40, Mar./Apr. 2011.

[43] J. Pukite and P. Pukite*, Modeling for Reliability Analysis*. Piscataway, NJ, USA: IEEE Press, 1998.

[44] M. Rausand and A. Høyland*, System Reliability Theory: Models, Statistical Methods, and Applications*, 2nd ed. Hoboken, NJ, USA: Wiley, 2004 [Online]. Available: http://www.ntnu.no/ross/srt

[45] N. Rickard and D. Young, Bandwidth Doesn't Matter; Availability Drives Enterprise Network Costs, Gartner, Inc., Jul. 2013, Tech. Rep.

[46] C. D. Rold and F. Ridder, Behind the Cloud: The Rise of Industrialized, Low-Cost IT Services, Gartner, Inc., Feb. 2011, Tech. Rep.

[47] Y. Sakamoto, M. Ishiguro, and G. Kitagawa*, Akaike Information Criterion Statistics*. Tokyo, Japan: KTK Scientific, 1986.

[48] M. Sallé and C. Bartolini, "Management by contract," in *Proc. IEEE/IFIP Network Operations and Management Symp. (NOMS 2004)*, 2004, vol. 1, pp. 787–800.

[49] G. Schick and R. Wolverton, "An analysis of competing software reliability models," *IEEE Trans. Softw. Eng.*, no. 2, pp. 104–120, 1978.

[50] B. Schroeder and G. Gibson, “A large-scale study of failures in high-performance computing systems,” *IEEE Trans. Depend. Secure Comput.*, vol. 7, no. 4, pp. 337–350, 2010.

[51] D. Scott, How to Assess Your IT Service Availability Levels, Gartner, Inc., Apr. 2009, Tech. Rep.

[52] A. Snow and G. Weckman, “What are the chances an availability SLA will be violated?,” in *Proc. 6th IEEE Int. Conf. Networking, 2007 (ICN’07)*, 2007, pp. 35–35.

[53] A. Snow, G. Weckman, and V. Gupta, “Meeting SLA availability guarantees through engineering margin,” in *Proc. 2010 9th Int. Conf. Networks (ICN)*, Apr. 2010, pp. 331–336.

[54] S. Taylor, D. Cannon, and D. Wheeldon, Service Operation (ITIL), The Stationery Office, TSO, 2007.

[55] S. Taylor, M. Iqbal, and M. Nieves, Service Strategy (ITIL), The Stationery Office, TSO, 2007.

[56] The Open Group, TOGAF Version 9—The Open Group Architecture Framework, 2009.

[57] M. Vineyard, K. Amoako-Gyampah, and J. Meredith, “Failure rate distributions for flexible manufacturing systems: An empirical study,” *Eur. J. Oper. Res.*, vol. 116, no. 1, pp. 139–155, 1999.

[58] N. Yang and B. Dhillon, “Availability analysis of a repairable standby human-machine system,” *Microelectron. Rel.*, vol. 35, no. 11, pp. 1401–1413, 1995.

Ulrik Franke received his Ph.D. in 2012, and his M.Sc. in 2007, both from the Royal Institute of Technology (KTH) in Stockholm.

He is a senior scientist at the Swedish Defence Research Agency (FOI). His research interests include Enterprise Architecture, high availability IT services, and the impact of ICT on politics and national security.

Hannes Holm received his Ph.D. in 2014 from the Royal Institute of Technology (KTH), and his M.Sc. in 2010 from Luleå University of Technology (LTU).

He is a scientist at the Swedish Defence Research Agency (FOI). His research interests include software security and the security of critical infrastructure control systems.

Johan König received his Ph.D. in 2014, and his M.Sc. in 2008, both from the Royal Institute of Technology (KTH).

He is an IT management consultant at Connecta AB. His research interests include quality analysis of active distribution grids from an ICT perspective.