Incident, Problem, & Change Management Metrics Benchmarks Update

A Report From The Pink Elephant IT Management Metrics Benchmark Service

Version   **: 1.0**
Date   **: July, 2012**

1   EXECUTIVE SUMMARY

The Pink Elephant IT Management Metrics Benchmark Service collects, analyzes and presents IT management metrics benchmarks. This Incident, Problem, and Change Management Metrics Benchmark update presents an analysis of voluntary survey responses by IT managers across the globe since early 2010. The surveys have thus far been limited to simpler metrics and the processes most broadly practiced.

Key points in this analysis:

Incident Management:

The number of Incidents is most influenced by (in order of influence)

1. The size of an IT organization (measured by quantity of IT Full Time Equivalent workers (FTEs) of all kinds (employees, contractors, and direct service providers' workers)
2. The number of users, and
3. The number of years that formal Incident Management has been in practice

At least a quarter of all respondents have no documented basis for any Incident Resolution Interval.

Problem Management:

The number of Problems added every month is just below the number of Active Problems already in progress (Problem Work In Progress (WIP)). This implies that the exit rate (rate at which Problems get resolved – or at least closed), must be pretty close to the number of new problems recorded every month. The 4.4 month average, problem average age at closure implies that some Problems are being closed very quickly – perhaps too quickly.

Change Management:

Among the several interesting metrics is an average 90% Change Executed Right First Time (no rollback or cancelation, and as scheduled). This appears to indicate that 10% of all Changes fail in at least 1 of the 3 ways, a quite disappointing benchmark.

Detailed observations follow:

This is an update on the Pink Elephant IT Management Metrics Benchmarks Service. Please submit your organization's data! The more participants in our on-line metrics benchmark surveys, the better! Surveys are available at https://www.pinkelephant.com/MetricsSurvey/.

We welcome your feedback. Please comment about this whitepaper on the blog post to let us know what you think, or write us at feedback@pinkelephant.com.
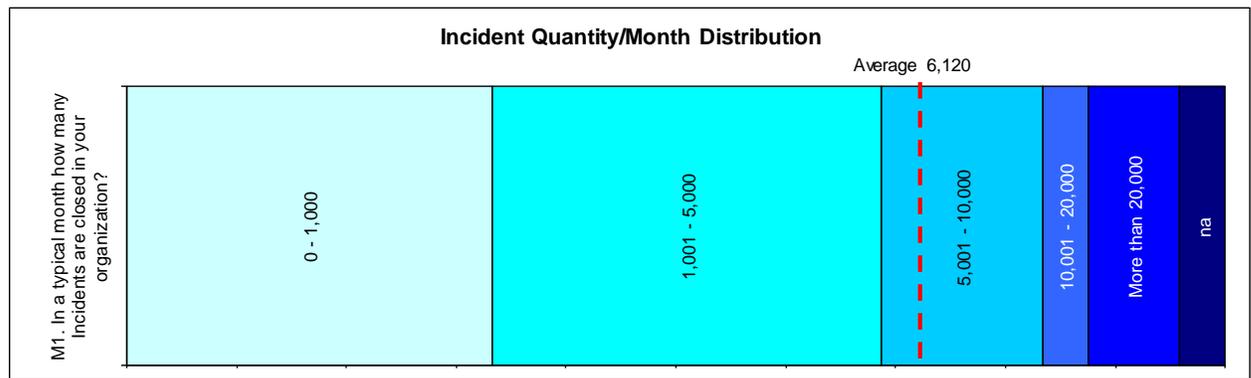
Table Of Contents

ITIL® is a Registered Trade Mark of the Cabinet Office.

2   INCIDENT MANAGEMENT

Incidents Per Month (Incident Rate):

| M1. In a typical month, how many Incidents are closed in your organization?					
0 - 1,000	1,001 - 5,000	5,001 - 10,000	10,001 - 20,000	More than 20,000	na
33%	35%	15%	4%	8%	4%

Incident Quantity/Month Distribution

Average 6,120

The data is widely distributed, and strongly favors lower rates. Not surprisingly, the size of the IT organization (measured by the number of IT FTEs) has the strongest relationship with Incident Rate: a positive correlation factor of .47.

First Contact Resolution (FCR Rate):

M2. Of the Incidents closed in a typical month, what percent are resolved by the first person contacted?					
Less than 70%	70% - 79%	80% - 89%	90% - 96%	97% - 100%	na
42%	22%	17%	9%	7%	3%

Incident First Call Resolution Distribution
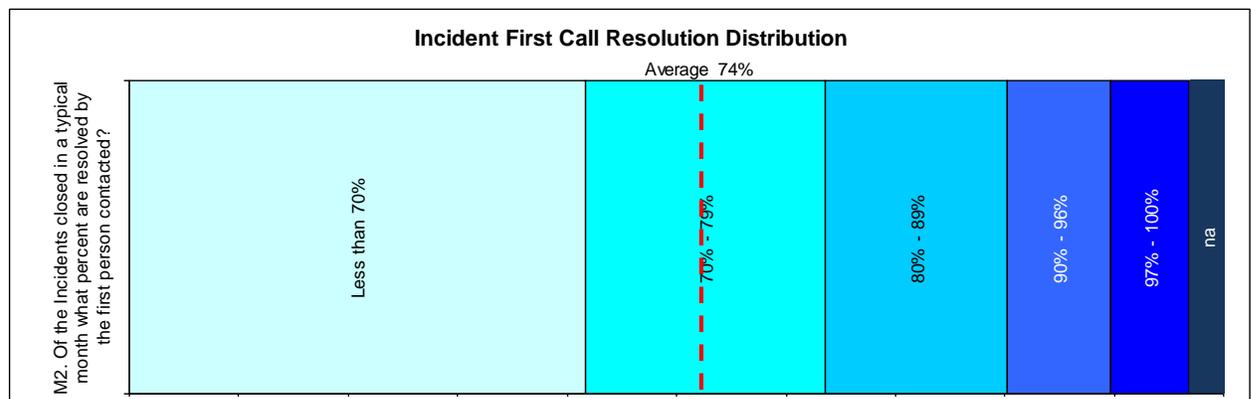
Average 74%

The average FCR rate is surprisingly low, and the distribution is strongly skewed to the low end. Also surprising is that the strongest correlation - a positive .33 with Incident Resolution Interval - is not stronger.
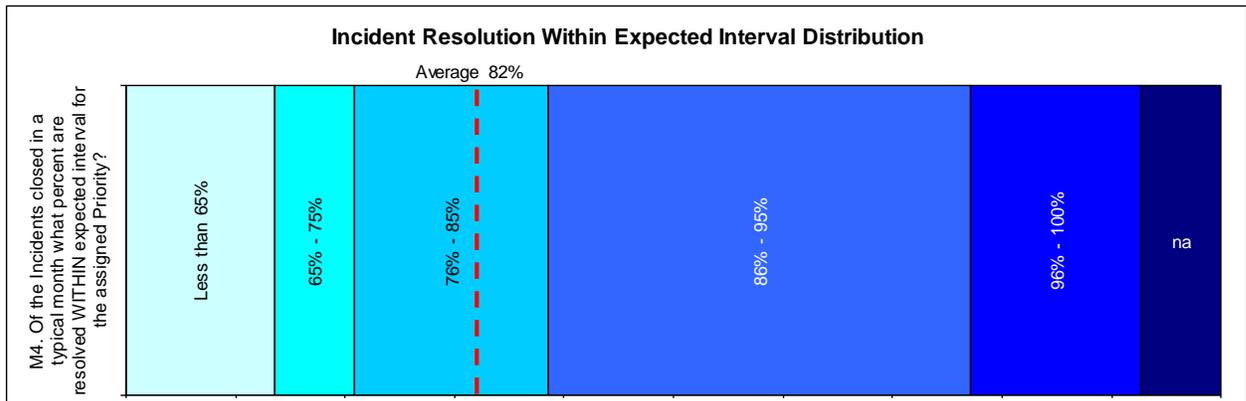
Incident Maximum Priority:

M3. Of the Incidents closed in a typical month, what percent are assigned the highest Priority?				
0% - 2%	3% - 7%	8% - 15%	More than 15%	na
40%	30%	14%	13%	4%

Incident Maximum Priority Distribution

Average 6%

| 0% - 2% | 3% - 7% | 8% - 15% | More than 15% | na |
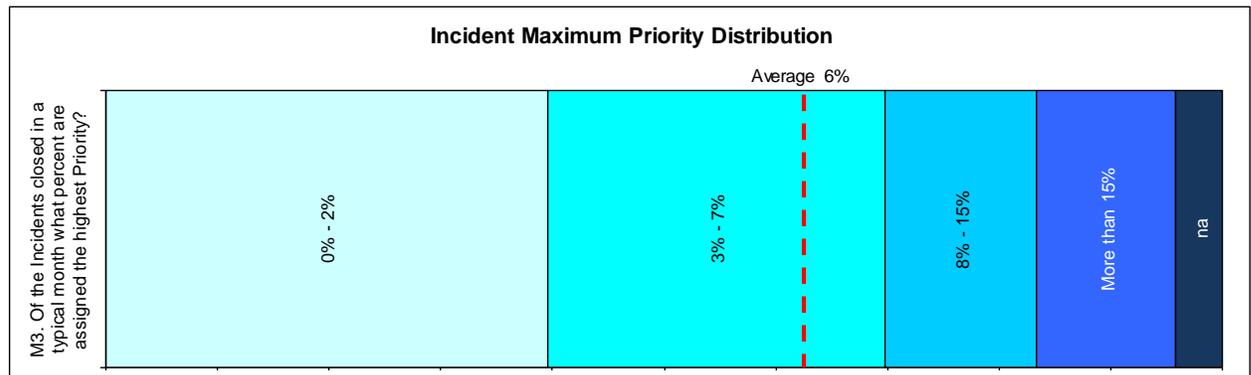
This appears to be a healthy sign that the "worst case scenarios" are fairly rare. 70% of respondents have fewer than 8% of the Incidents pegged at the maximum Priority. As you might expect, the strongest correlation is an inverse relationship (-.33) with the Incident Resolution Interval. Organizations with low On-time Incident Resolution Rate (longer running Incidents) also have more Incidents getting the highest priority. A vicious cycle.

Incident Resolution Within Expected Interval:

M4. Of the Incidents closed in a typical month, what percent are resolved WITHIN the expected interval for the assigned Priority?					
Less than 65%	65% - 75%	76% - 85%	86% - 95%	96% - 100%	na
14%	7%	18%	39%	16%	7%

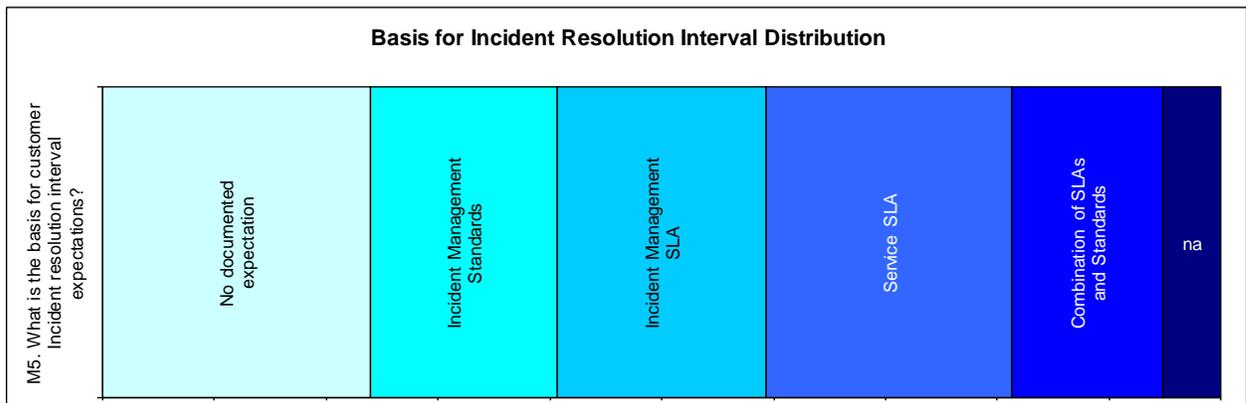**Incident Resolution Within Expected Interval Distribution**

Average 82%

M4. Of the Incidents closed in a typical month what percent are resolved WITHIN expected interval for the assigned Priority?

| Less than 65% | 65% - 75% | 76% - 85% | 86% - 95% | 96% - 100% | na |

There is a lot of room for improvement here. This metric has a positive .33 correlation with First Contact Resolution and a negative -.33 correlation with the percentage of Maximum Priority Incidents. It could be that a poorly defined Incident Resolution Interval Expectation may an underlying cause (see next item).

Basis For Incident Resolution Interval:

M5. What is the basis for customer Incident Resolution interval expectations?					
No documented expectation	Incident Management Standards	Incident Management SLA	Service SLA	Combination of SLAs and Standards	na
24%	17%	19%	22%	14%	5%

Basis for Incident Resolution Interval Distribution

M5. What is the basis for customer Incident resolution interval expectations ?

| No documented expectation | Incident Management Standards | Incident Management SLA | Service SLA | Combination of SLAs and Standards | na |

The most distressing information is almost a quarter of the respondents have NO basis for Incident Resolution Interval expectations. This is a tricky area. Service Level Agreements (SLAs) are a good basis, but can become difficult to comply with if there are many different services with varying and "special" requirements to track. Incident Management SLAs are OK – except that it is not good practice to set up Processes as Services. My personal favorite is Incident Management Standards. Each Incident Priority gets a Notice to Response and a Notice to Resolution Target Interval. SLA's availability requirements can be calculated and then Incident Priorities preset for outages.

It is interesting that organizations that have no documented Incident Resolution expectation span the entire field of survey respondents, with more than half being of medium size 50-750 IT FTEs and 100-10,000 users, and just under half are quite young or early in their journey (Under 2 years since Incident Management was implemented).

The impact of experience with a process and continual improvement:

The breadth and depth of the Incident Management data provides the following Incident Management metrics averages for "years of experience" (which also indicates the minimum age of the IT organization). We should soon have the same breadth and depth for the Problem and Change metrics:

A1. How long ago was Incident Management first implemented in your organization?

| Metrics Question	Less than 2 years ago	2 - 4 years ago	4 - 6 years ago	6 - 10 years ago	More than 10 years ago
M1. In a typical month, how many Incidents are closed in your organization?	3,391	5,479	6,474	9,500	12,438
M2. Of the Incidents closed in a typical month, what percent are resolved by the first person contacted?	76.5%	72.8%	75.9%	64.2%	78.3%
M3. Of the Incidents closed in a typical month, what percent are assigned the highest Priority?	8%	4%	7%	5%	5%
M4. Of the Incidents closed in a typical month, what percent are resolved WITHIN the expected interval for the assigned Priority?	81%	83%	80%	90%	74%
M5. What is the basis for customer Incident Resolution interval expectations?	IM Std.	IM SLA	IM SLA	IM SLA	IM SLA

The responses by the number of users and/or IT FTEs (IT organization size indicator) have been less broad. As the population of the database grows, reliable implications of these factors will be analyzed and provided.

3   PROBLEM MANAGEMENT

Problems Per Month (Problem Rate):

| M1. In a typical month, how many new Problems are recorded in your organization?					
0 - 10	11 - 50	51 - 100	101 - 200	More than 200	na
29%	50%	4%	8%	4%	4%
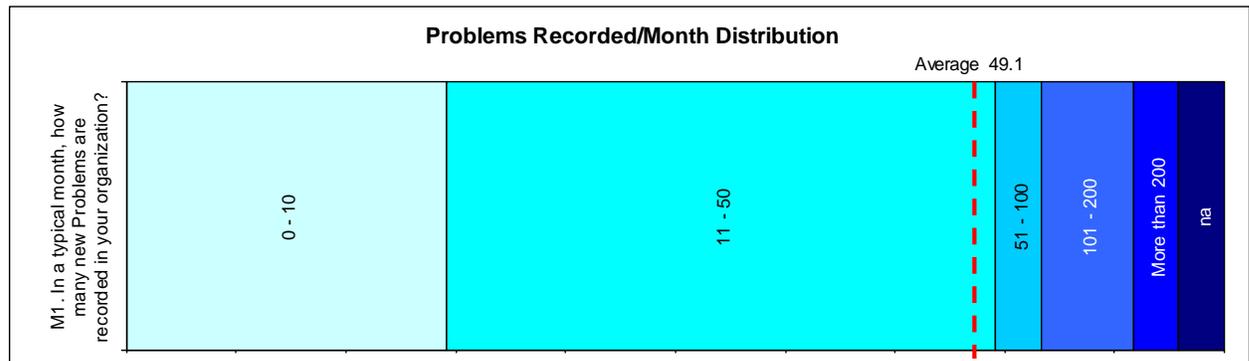
The Problem Addition Rate is less than 1% of the Incident Rate. This indicates that – as recommended – most organizations' Incidents have a strong Pareto bias, and the organizations open Problems for only the most important subset of Incidents – and then take a while to close them (see Age at Closure below). It is interesting that the strongest covariance for this metric is with Problem WIP (.58 positive correlation) and the Number of Users (.49 positive correlation). The correlation with WIP is understandable but worrisome since the Problem Rate exceeds WIP. If this keeps up, WIP will continue to grow – not a good thing – unless the exit (closure) rate can be improved or the Problem initiation threshold criteria increased.

The correlation with the number of users may be because the number of users may drive a larger set of discrete systems and applications – and therefore more discrete points of failure or unique errors at the root of more problems.

Problems & Known Errors In WIP:

M2. How many Problems and/or Known Errors are actively being worked on (researched, solutions being determined, mitigated/eliminated)?				
0 - 15	16 - 30	31 - 60	61 - 200	na
38%	17%	13%	25%	8%

Problems & Known Errors WIP Distribution

Average 48.8
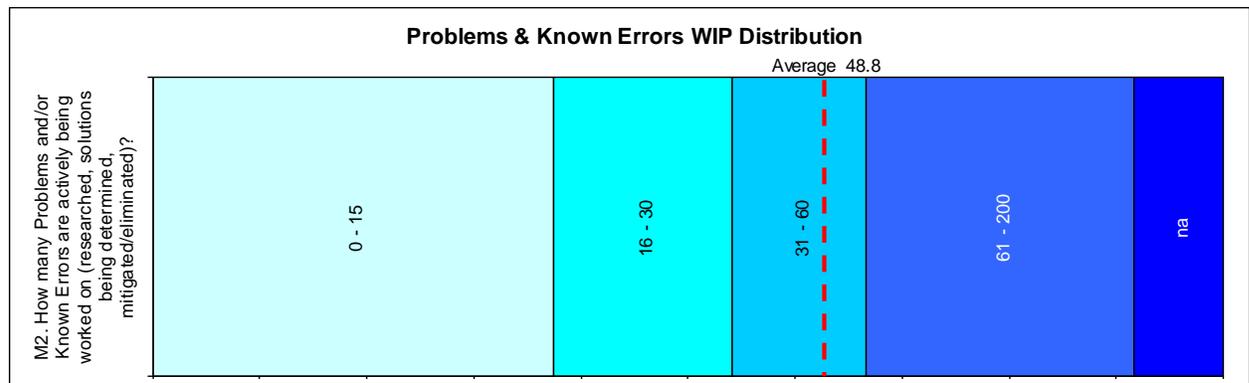
As noted above the Problem Rate and WIP are very close and enjoy a .58 positive correlation factor. This implies that the Problem closure rate is very near the Problem rate – if the WIP quantity is stable. But, with the average age at closure being 4.4 months there are some very different numbers hidden in the "not available" responses. Time will tell.

Problems Assigned Highest Priority:

M3. Of the Problems and Known Errors being worked on, what percent of them are assigned the highest Priority?				
0% - 2%	3% - 7%	8% - 15%	More than 15%	na
33%	21%	13%	17%	17%

**Problems Assigned Highest Priority Distribution**
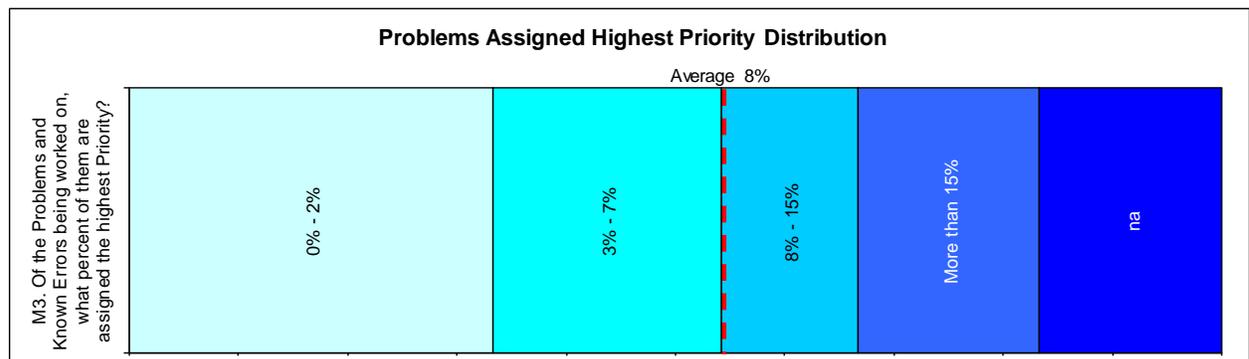
Average 8%

At 8% of Problems at highest priority, this also compares closely with Incident Management's 6% Maximum Priority average. The strongest positive correlations are with the Number of Users and New Problems per month per user and per FTE. The strongest negative correlations are with the Number of Years since Problem Management was implemented, and the size of the IT organization (measured in IT FTEs). This metric may also indicate that Problem Management is

being practiced like Risk Management – which is a good model for Problem Management. In Risk Management, the priority is a function of the Risk Exposure: the business impact in lost business cost and recovery cost per occurrence times the probable number of occurrences per year.

Problem / Known Error Age At Closure (Months):

M4. Over the last year, what has been the average age of Problems and/or Known Errors at closure (Closure: customer accepts implemented resolution or mitigation, or customer accepts risk)?				
1 month	2 - 3 months	4 - 6 months	7 - 12 months	na
4%	38%	29%	13%	17%
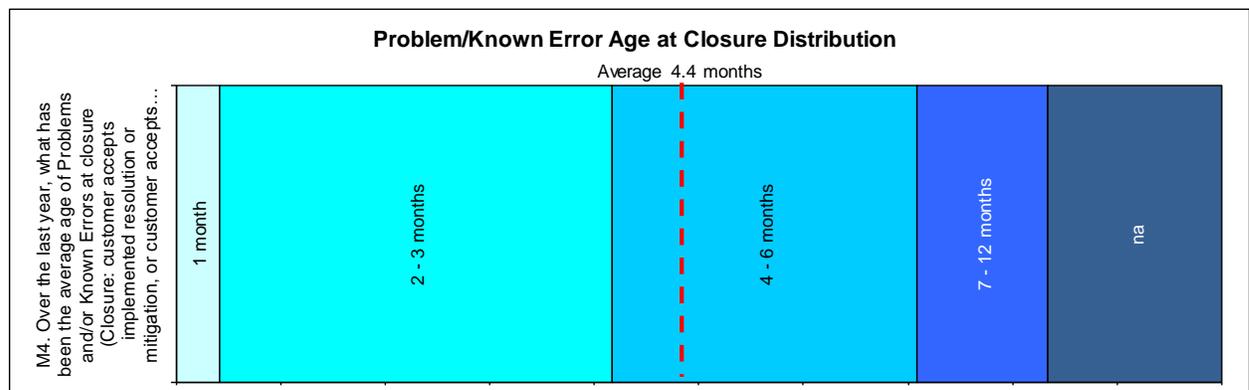
Not knowing what policies govern the closure of Problems, it is difficult to assess the roots of this metric. The only strong correlation is the WIP / User (positive .45), indicating that spreading the pain/risk across users helps build a stronger case to close Problems – again, depending on the closure requirements.

Problems Without Known Errors:

M5.What percent of the current Problems and Known Errors are Problems that have no Known Error? (Current Problems without Known Errors divided by all Current Problems and Known Errors)					
Less than 20%	20% - 39%	40% - 59%	60% - 79%	80% - 100%	na
21%	21%	25%	4%	13%	17%

Problems wihout Known Error Distribution

Average 42%

M5.What percent of the current Problems and Known Errors are Problems that have no Known Error? (Current Problems without Known Errors divided by all Current Problems and Known Errors)

| Less than 20% | 20% - 39% | 40% - 59% | 60% - 79% | 80% - 100% | na |
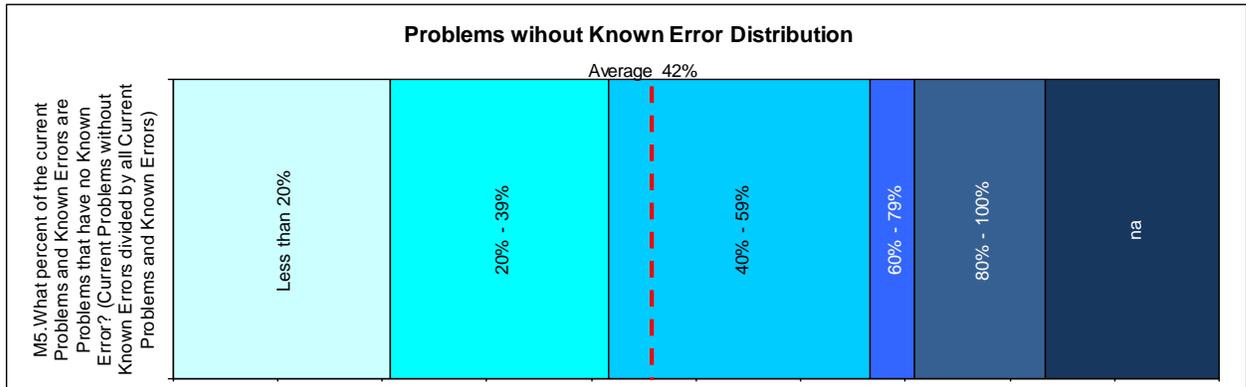
If 42% of Problems have no Known Error, then almost half of the Problems are still in the Root Cause determination phase. It cannot be determined from the survey why this is the case. It is somewhat comforting to imagine that the Problems with Known Errors are in active mitigation, but they are more likely waiting their turn in the project portfolio queue. Only WIP correlates with this metric at positive .33. The more slowly (and poorly) Known Errors are identified, the more the Problem WIP will grow.

4 CHANGE MANAGEMENT
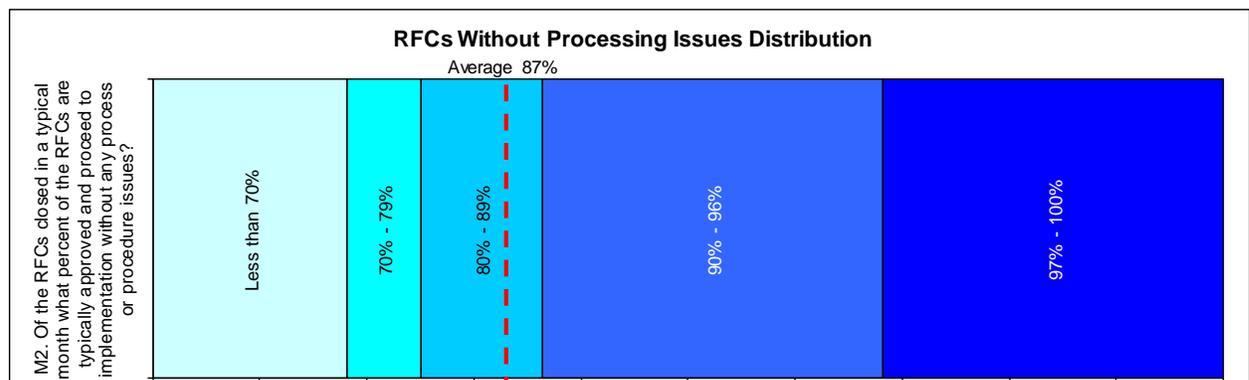
Requests For Change (RFCs) Closed Per Month (RFC Rate):

M1. In a typical month, how many Requests for Change (RFCs) are closed in your organization?						
0 - 50	51 - 100	101 - 500	501 - 1,000	1,001 - 2,500	2,501 - 5,000	More than 5,000
18%	11%	34%	14%	5%	9%	9%
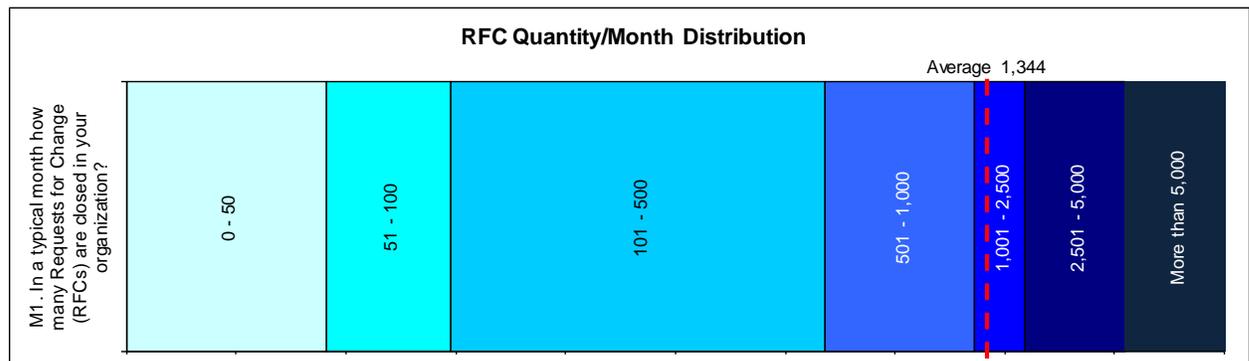
Average monthly RFC rates this high are an indicator of the risks in IT Operations environments. It is no surprise that the RFC rate has a very high positive correlation (.67) with the quantity of IT FTEs – a proxy measure for the size of the IT organization) and, of course a positive .75 correlation with RFC/FTE. Since the surveys are discrete by process, there is no data on the correlation of Incident, Problem, and Change rates.

RFCs With No Processing Issues (RFC Approval Right First Time Rate):

M2. Of the RFCs closed in a typical month, what percent of the RFCs are typically approved and proceed to implementation without any process or procedure issues?				
Less than 70%	70% - 79%	80% - 89%	90% - 96%	97% - 100%
18%	7%	11%	32%	32%

There is an ITIL legend that the RFCs approval and processing is a trap that slows the rate of change. This 87% RFC Approval Right First Time Rate undermines the legend. Almost 30% have RFC Approval Right First Time Rates above 97%. There is a very strong positive correlation with the RFC Execution Right First Time Rate, indicating that IT organizations with efficient RFC processing procedures also have higher Change Execution Right First Time rates.

Looking on the dark side, all of this could also indicate weak RFC controls. That will be the subject of further study.

RFC Right First Time (Change Execution Right First Time Rate):

| M3. Of the RFCs closed in a typical month, what percent of the RFCs are typically implemented successfully as scheduled the first time?			
Less than 80%	80% - 90%	90% - 96%	97% - 100%
18%	16%	30%	36%
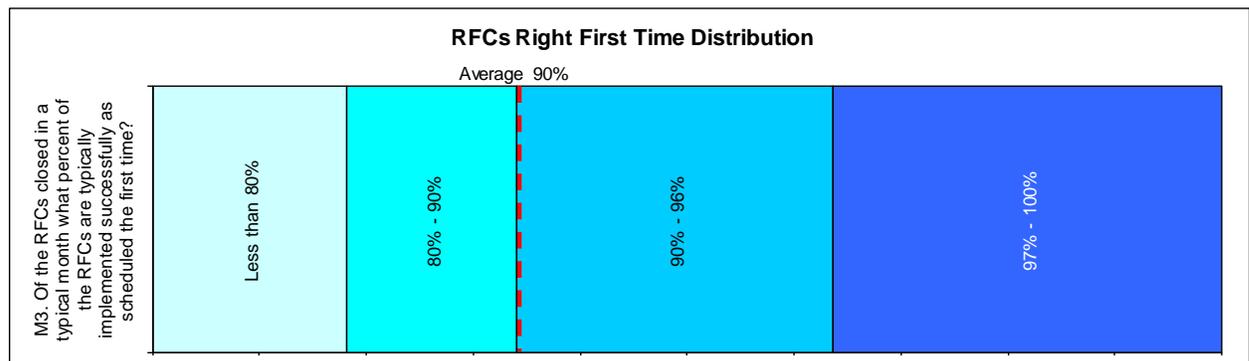
A Change Right First Time (RFT) rate of 90% is fair, especially since it includes executing on time as well as without having to roll back the change. Still, a 10% failure rate seems very high. This metric has three strong correlations.
- As mentioned above, there is a positive .67 correlations with RFC Approval Right First Time Rates
- There is a negative -.35 and .-34 correlation respectively with the Emergency RFC Rate and the number of days between RFC submission and closure, grounding the often repeated warning that high pressure (Emergency) and haste make waste

Emergency RFC Rate:

M4. Of the RFCs closed in a typical month, what proportion of the RFCs are Emergency Changes?					
0% - 1%	2% - 4%	5% - 9%	10% - 15%	More than 15%	na
9%	27%	18%	25%	18%	2%

Emergency RFC % Distribution

Average 9%

M4. Of the RFCs closed in a typical month what proportion of the RFCs are Emergency Changes?

| 0% - 1% | 2% - 4% | 5% - 9% | 10% - 15% | More than 15% | na |
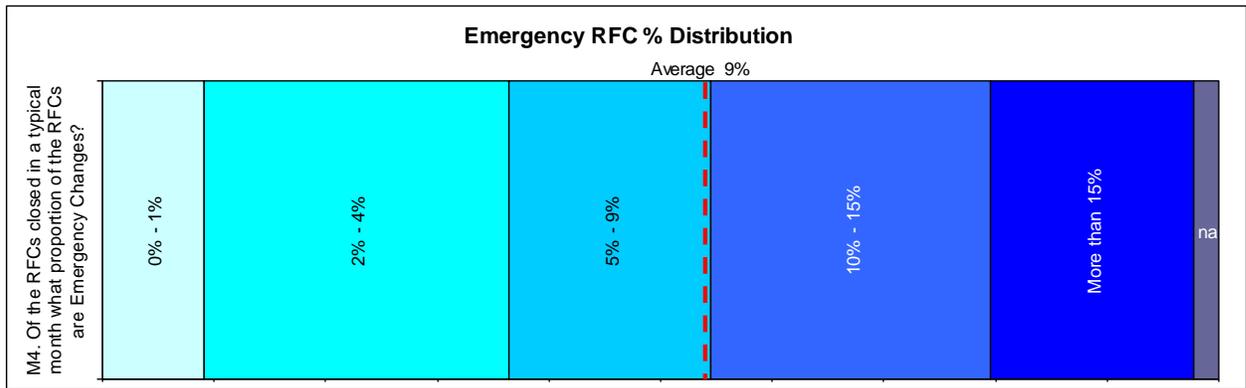
It is not possible to determine for certain why the Emergency RFC rate is this high for almost half of the survey respondents. It could be high priority Incidents (6% of all Incidents in our IM benchmark), or "Emergency" project work. (Failure to plan or manage performance on a Project Team's part IS an "Emergency" on Change Management's part.) The strongest correlations are the negative -.35 with Change RFT (above), and positive .30 correlation with RFCs per User.

RFC Life Expectancy (Days From RFC Submission To Closure):

M5.What is the average interval between RFC submission and RFC closure?						
1 - 5 days	5 - 15 days	15 - 30 days	30 - 60 days	60 - 90 days	More than 90 days	na
7%	32%	25%	16%	11%	5%	5%

RFC Life Expectancy Distribution

Average 32 days

M5.What is the average interval between RFC submission and RFC closure?

| 1 - 5 days | 5 - 15 days | 15 - 30 days | 30 - 60 days | 60 - 90 days | na |
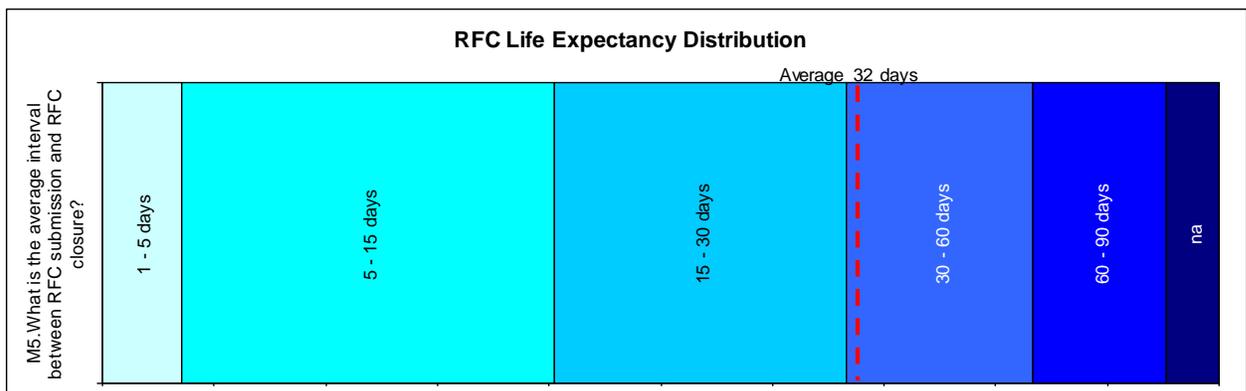
A 32 day average seems a long time for the average RFC to be open, especially with almost 40% indicating 1-15 days. This may be the result of a large variance in Change scope across the responding organizations, with program level changes open for a very long time at the high end of the scale. The negative -.34 correlation with Change Right First Time affirms that the metrics are at least directionally correlated. The mildly negative -.24 correlation with Years since Change Management Implementation indicates that once one has been practicing Change Management for a long time, the interval between submission and closure may shorten.