Process improvement focused analysis of VINST IT support logs

Shobana Radhakrishnan and Gokul Anantha

SolutioNXT Inc.

(www.solutionxt.com)

Abstract

The goal of this paper is to use existing logs and transactional information from a Support and Problem Management system called VINST, perform detailed analysis of various efficiency and performance factors and identify some key actionable patterns for improvement. We have used a combination of process discovery tools (such as Disco) and reusable scripting on MS Excel to perform this analysis. The focus of our approach is to discern findings and encapsulate them within real world perspectives. We brought this real world perspective in reclassifying the given dataset into a) All cases b) Incidents only b) Incidents escalated to problems and c) Problems only. We assessed a) wait status abuse, b) ping –pong behavior across levels and across teams and c) general case flow pattern. We uncovered interesting finding and captured a set of clear recommendations based on these findings.

Overview

We received three sets (or files) of logs (incidents, open and closed problems) containing data set for the period 11.01.2006 to 15.06.2012. The logs contain the following information on all service requests (SR's) - Creation Date, Status, Sub-Status, Support Team & org, level of Impact, product, country and support owner. Detailed information to understand the process, terminology and support ticket workflows was provided in the VINST manual and the document 'description of the dataset and questions '.

As practitioners, we have taken a 'matter–of- fact' approach to analysis, comprising of the following steps:

1. Understand the contextual nature of incident and problem management at Volvo, Belgium
2. Dissect the transaction logs received to determine patterns and answers for key questions raised in the challenge. Also provide our observations on general patterns.
3. Overlay our domain knowledge and understanding of IT support systems to arrive at recommendations

The succeeding sections detail out each of these steps.

Volvo IT Support system

- Volvo Belgium's IT support system comprises of three levels
    - First line = service desk or expert desk – comprising of service desk front desk, offline desk or desk-side support. These are local and global teams not within an Org line.
    - Second line comprises of specialized functional teams within an org line (for example: Org line C or Org line A2).
    - Third line is a team of specific product or technical experts and is also within an Org line.

  A complete list of Org lines, functional teams and the associated support teams and as deciphered from the available data is presented in Appendix 1
- All SR's are classified and prioritized [1]based on a matrix rule set based on impact and urgency. Impacts is correlated to # of people/systems affected and there are four impact levels viz. major, high, medium and Low. Urgency determines the required speed of solving and has three levels viz. high, medium and low. The SR creator (user) has influence on urgency at the time of incident creation. Also, while SR's can be upgraded in terms of emergency (& this may have impact on case routing), impact definitions cannot be upgraded. A major impact SR[2] is attached highest priority and SLA norms do not apply.
- Problem Management is the process of managing escalated incidents, 'major' impact incidents and root cause analysis (RCA) for 'complete'[3] incidents. Problem has four stages viz. Queued, Accepted (Assigned, Awaiting Assignment, Cancelled, Closed, In-Progress, Wait, Unmatched), Completed, Closed

Analysis: Validating existing datasets

1. There are a total of 9395 unique SR's across the 3 log files, broken down as below
    a. 'Incident' file = 7554 SR's
    b. 'Closed Problems' file = 1487 SR's,
    c. 'Open Problems' file =819 SR's.

   There are no duplicate SR's between a) the 'Incidents' file and 'Open Problems' file and b) 'Incidents' file and 'Closed Problems' file. However, 465 SR's overlap 'Open Problems' and 'Closed Problems' files.

---

[1] This is analogous to a severity definition commonly used in many other organizations.

[2] Also sometimes identified as a severity 1 incident in other organizations.

[3] Have used 'complete' instead of a generic term 'closed' to reflect typical assigned status of such SR's in Volvo Belgium.

2. Based on column 'Involved ST', we were able to clean up and separately capture the team level handling a particular SR transaction and stored it in a new column 'ST level'[4]. For example, if an SR was handled by support team 'G199 3rd', we have classified it as a level 3 case for that transaction. There are some cases where the data is shown with multiple teams example "V13 2nd 3rd" for SR 1-364285768. In such circumstance, we have classified that transaction as being handled by level 2 ST. There are only a few such SR instances, so we do not expect any impact from this approach.
3. Classification of SR transactions by support team level provides us a view of the SR flow across levels. This view allows us to re-examine the SR distribution between the three files. On closer examination of the 'Incidents' file, for example, we determined that SR transactions were not limited to any level in the file. The distribution of SR transactions across ST levels in each of the file is captured below in Table 1.

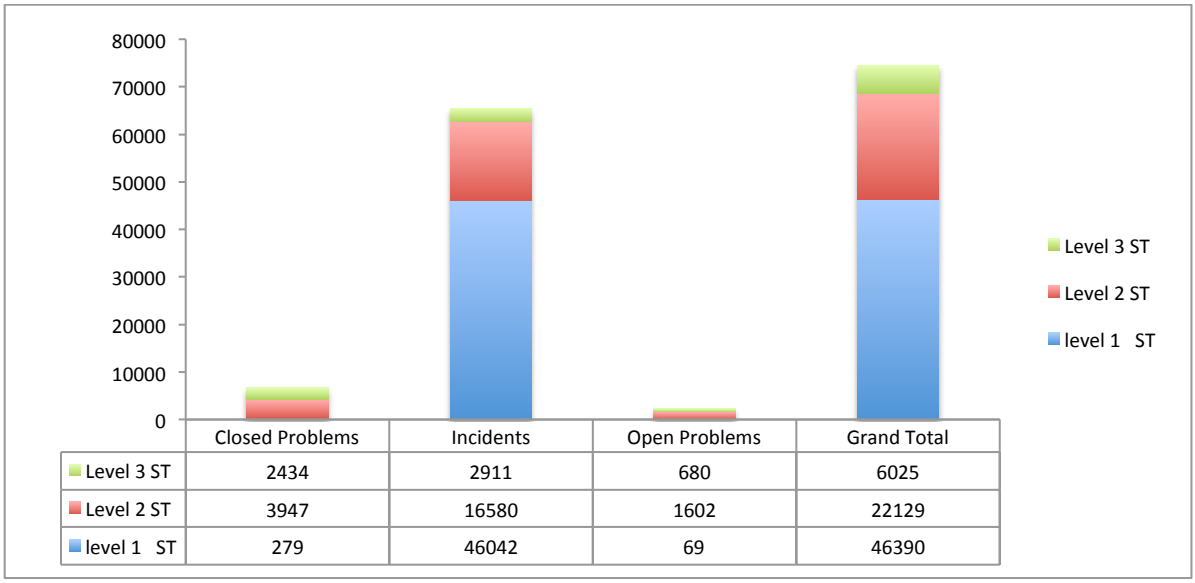

	Closed Problems	Incidents	Open Problems	Grand Total
Level 3 ST	2434	2911	680	6025
Level 2 ST	3947	16580	1602	22129
level 1 ST	279	46042	69	46390

Table 1: Distribution of SR transactions across levels

- ✓ 19,491 of 65,333 (Approximately 30%) SR transactions in 'Incident file' were handled by level 2 and level 3 teams. Perhaps case priority is resulting in their routing to the higher levels (and we will analyze this in the next section).
- ✓ ~ 4% of the total transactions in Open and Closed Problems files were handled by level 1 teams. We would have expected to see a higher number, but it is likely that the cases were identified as "problems" before being created and directly assigned to the concerned teams. We will present a more accurate assessment of this in the subsequent sections.

[4] The updated dataset is provided in appendix 1

- More importantly, this information highlights that the individual files cannot be used independently for analysis. We examination the merged dataset (using Disco) with 'activity' set as a transition from one ST level to another, and this further validated our assumption. The outcome (fig 1 below) showed cased directly being assigned to ST level 1, 2 and 3 and then flowing between these teams.

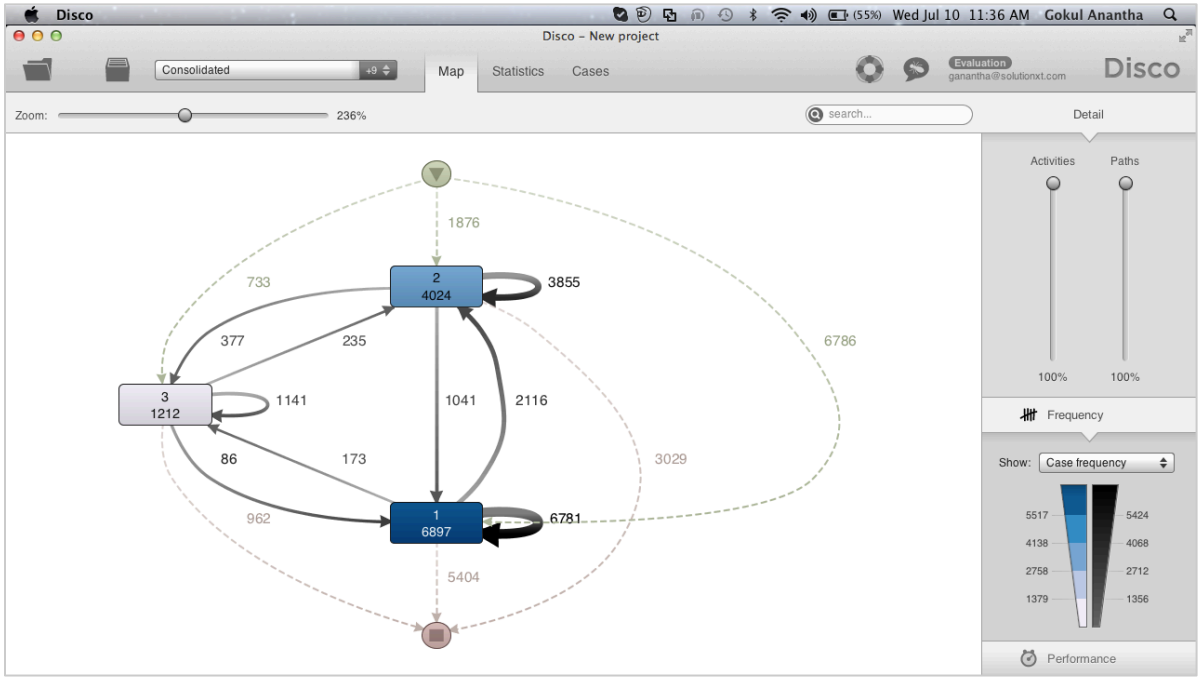

Fig 1: flow analysis of SR transaction across ST levels.

4. As a consequence of above assessment, we merged all the SR's and their transactions into a single 'master' file while retaining a reference for the origin file. To re-create datasets for analysis, we went back to first principles
    a. Incident – A case that is handled in its entirety by a first level team. By definition, not all SR's are incidents.
    b. Problem – An escalated SR or an SR that requires specialized handling. A problem could be a defect (an SR that might need a quick fix or a work-around) or an enhancement(SR's with major impact that need technical fixes delivered either as a patch or as an application release)
5. Based on above, we re-classified the merged 'master' dataset into the following for detailed analysis.
    a. Incidents only
        i. SR's with all transactions handled by level 1 support teams only
    b. Problems only

i. SR's with all transactions handled by level 2 or level 3 support teams and none handled by a level 1 support team because of being directly assigned to a level 2 ST or level 3 ST.

c. Incidents that escalated to problems

i. SR's that were escalated up by a level 1 ST or ones that were pushed down to a level 1 support team for re-assignment to a different level 2 /3 ST. An example SR is 1-475885658 which is routed to all levels in its 29 transaction journey over 1 year and 68 days (please see flow in fig 2 below).

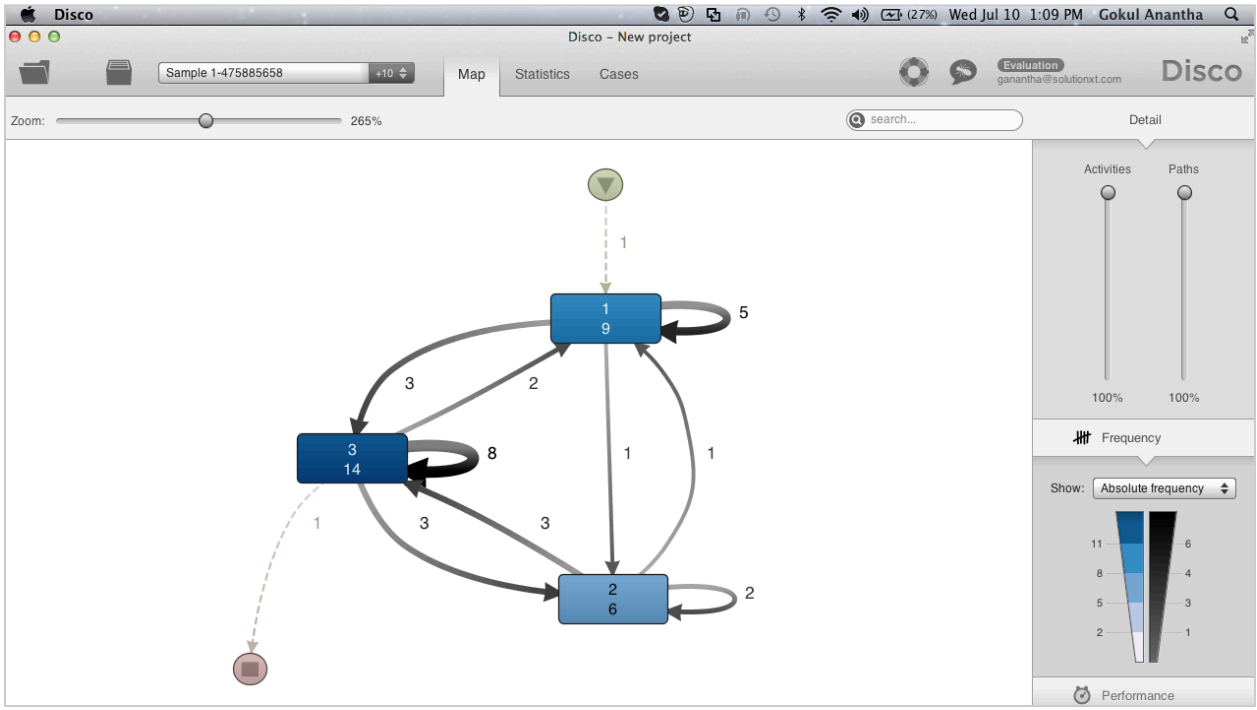

Fig 2: SR example with transition across all ST levels

6. To allow automated handling, the rule governing which dataset an SR is bucketed into is illustrated below:

SR Transaction handled by			
Level 1	Level 2	Level 3	SR Dataset
Yes	No	No	Incidents only
No	Yes	Yes	Problem only
No	No	Yes	Problem only
Yes	Yes	Yes	Incidents escalated to problems

Yes	No	Yes	Incidents escalated to problems
Yes	Yes	No	Incidents escalated to problems
No	Yes	No	Problem only

The next section details out our analysis of the master data set and subsets as listed above. This classification was particularly beneficial in identifying patterns such as Push to Front and Ping Pong, as well as make additional observations related to other parameters.

Analysis: Identifying patterns

1. Analysis of all SR's (merged dataset)

- The merged dataset comprises of 9395 SR's and 75444 transactions.[5] A flow analysis of all SR's (using Disco) is provided in fig 3 below

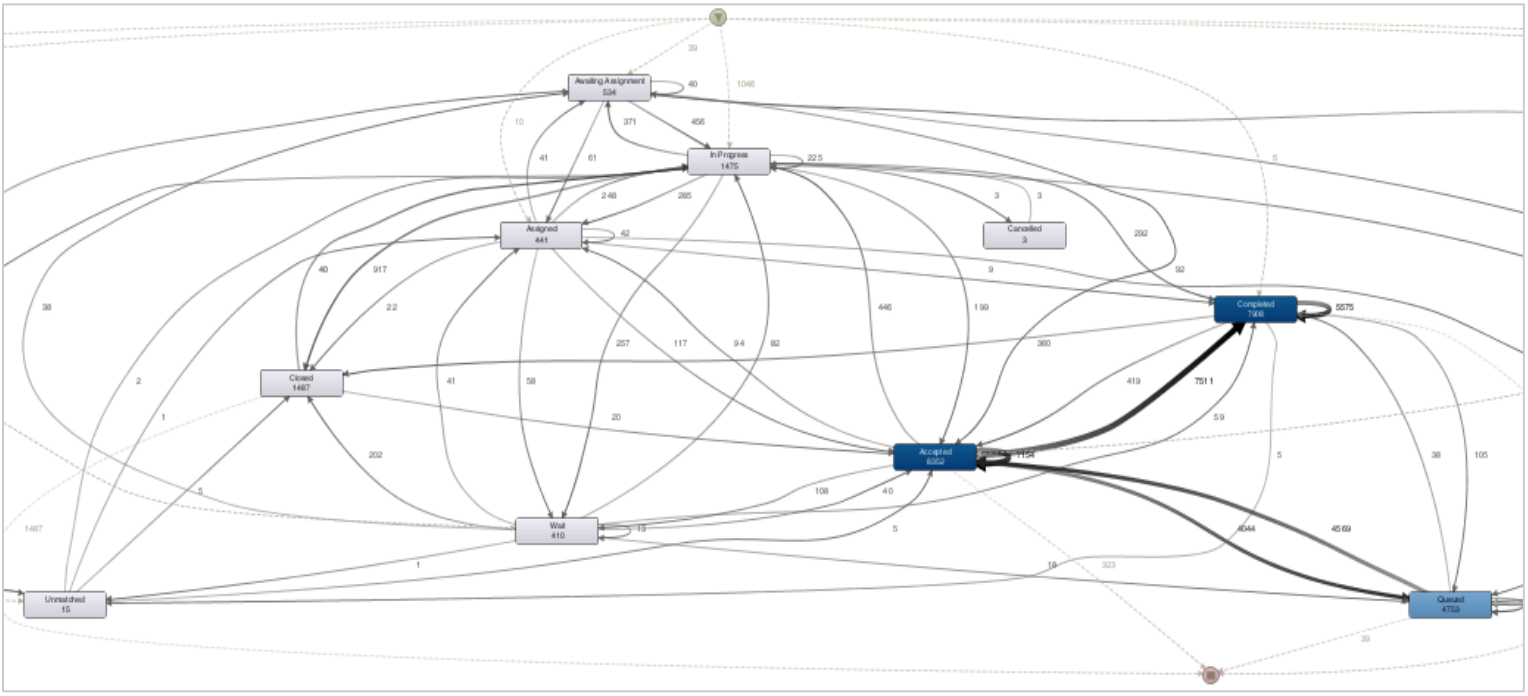

fig 3: SR process flow model (using change in 'status' as activity)

- Based on industry heuristics, we expected a state transition between statuses consisting of well-defined escalation paths as depicted in fig 4 below. Such a model also lends itself to easier analysis for continuous improvement.

[5] Available for reference in appendix 1

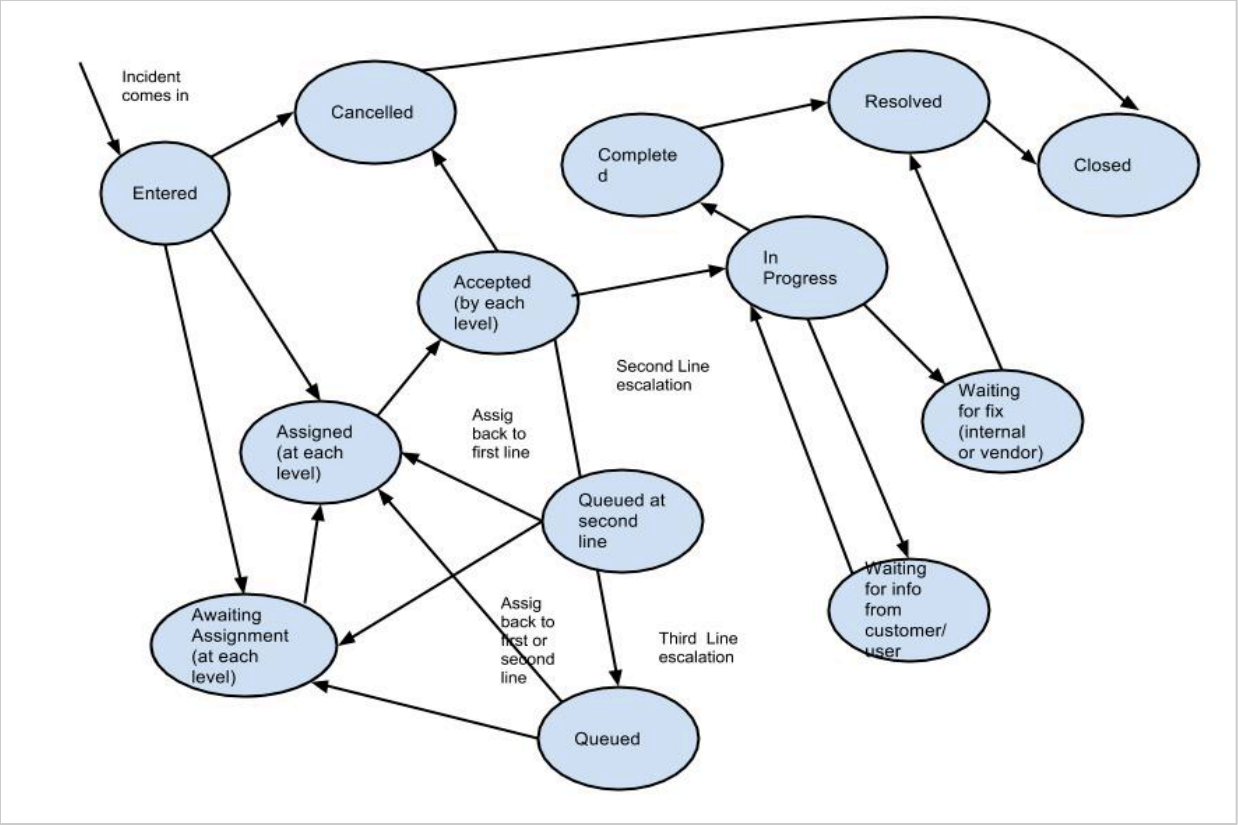

fig 4: Expected SR state transition (using Volvo IT statuses)

- As can be determined from fig 3 above, we observe a different model in effect. For example, 7084 SR's entered into "accepted" as their status at first entry instead of queued (which was the first entry status only for 1190 of the SR's analyzed). Also 1046 support requests were directly in the "In Progress" state at first entry instead of going through a queued state. Summary of firs entry state distribution below:
    - In Progress – 1046 times
    - Queued – 1190 times
    - Accepted – 7084 times
    - Other statuses (Completed, Awaiting Assignment, Wait, Unmatched) :75

  SR's being directly worked upon ("accepted" status) before triage ("queued" or 'assigned: status) could result in ping-pong behavior as the initial assignee may not be the right person/team to pick this up, resulting in increased ETA for resolution.
- Approximately 1882 net SR's are in 'In Call' status. It is not clear what the status means, but assuming this indicates resolution during initial phone contact (i.e. without the prior need for an SR ticket), this can become the highest efficiency level organizations can aspire to achieve. If this is

intended to capture the means of contact, it is better reflected as a channel, rather than a status. Also, if this channel proves to be an increasingly frequent source, then the "push to front" approach should focus on a first contact resolution (FCR) metric. There are structural implications in moving to such a metric and we will discuss this in the recommendations section

- We also found that a few status and sub-status values did not have clear demarcation of usage. For example a status of "closed" vs. a sub-status of "completed" and vice-versa, likewise a status of "accepted" and a sub-status of "assigned" and vice-versa seem to imply the same state and can potentially be cleaned up (please see table 2 below). This kind of cleanup can facilitate cleaner progression analysis.

	Sub-Status					
Status	Cancelled	Closed	Completed	In Call	Resolved	Grand Total
Closed			1565			1565
Completed	1	6103		2035	6115	14254

Table 2: Status Vs. Sub-status (overlapping usage?)

- Wait status usage

Is there evidence of wait status abuse?

- 'Wait' as a status is only used 410 times. So, prima-facie, it appeared not.
- However, a deeper examination provides a different and revealing perspective. We analyzed sub-statuses most used along with the various statuses. We found that the various 'Wait-xx' sub- statuses were actually used in conjunction with 'Accepted' status nearly 100% of the time (& not with a Wait status as we might have intuitively assumed). 6869 of 41698 transactions (16.5%) with initial status as 'Accepted' had some sub-status indicating wait (e.g. Wait-customer, Wait-implementation etc.).

	Sub Status																
Status	Accepted	Assigned	Awaiting Assignment	Cancelled	Closed	Completed	In Call	In Progress	Queued	Resolved	Unmatched	Wait	Wait - Customer	Wait - Implementation	Wait - User	Wait - Vendor	Grand Total
Accepted		3436						31393				1745	101	493	4217	313	41698
Assigned	614																614
Awaiting Assignment									875								875
Cancelled						3											3
Closed						1565											1565
Completed				1	6103		2035			6115							14254
In Progress	3066																3066
Queued			11927														11927
Unmatched											15						15
Wait	527																527
Grand Total	4207	3436	11927	1	6103	1568	2035	31393	875	6115	15	1745	101	493	4217	313	74544

Table: 3: Status vs. Sub-status overlay with emphasis on 'Wait-xx' sub-statuses

To explore further, we decided to explore the SR flow pattern using sub-status as the SR transition activity. Again, disco proved to be a pretty valuable tool for this purpose.

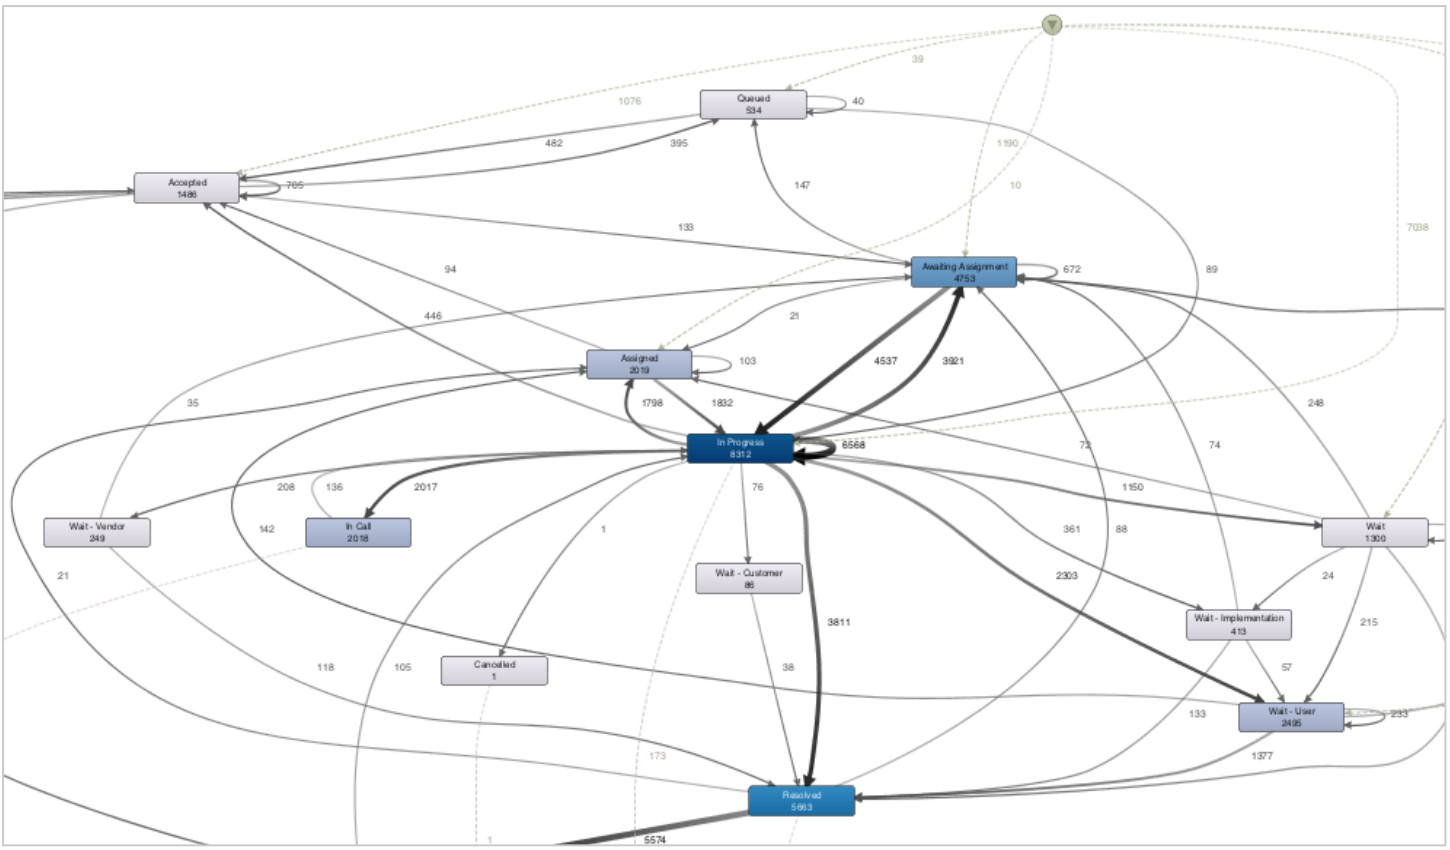

Fig 4: All SR's transaction flow by sub-status

Approximately 4367 SR's had a 'wait-xx' sub-status immediately following and in-progress' sub-status. This is 47% of all SR's.

- Next we created a smaller dataset of all unique SR's with at-least 1 status as 'accepted' and having at-least one 'wait-xxx' sub-status in its transaction logs. We found 3550 such SR's. When this dataset was analyzed at ST level, 3069 of the 3551 SR's were handled by a level 1 support team. This gave us a 'Immediate wait-usage ' index of 0.86 for level 1 ST . This analysis was further corroborated by another statistic as below

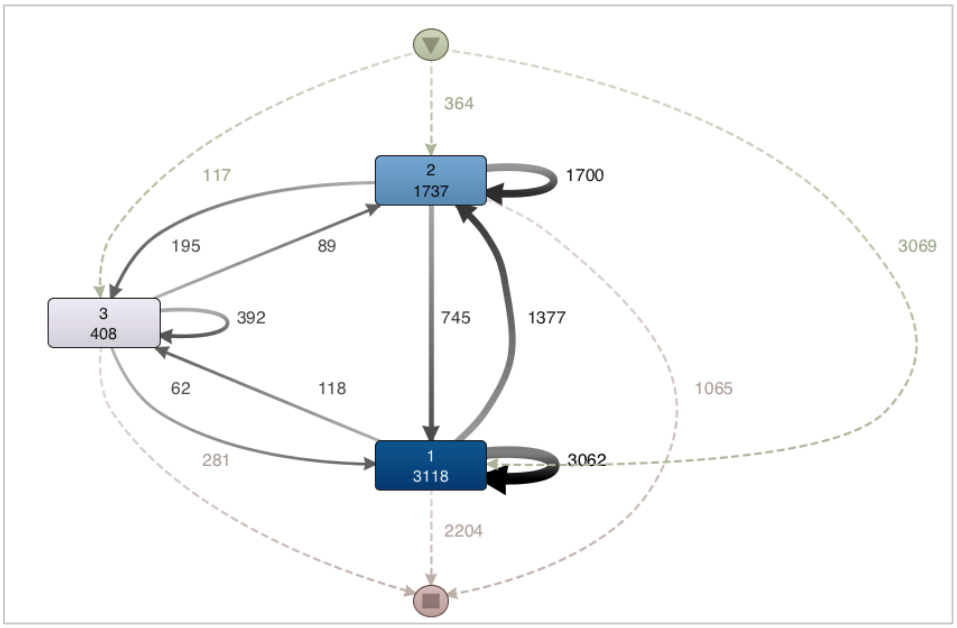

Fig 5: 'Wait-xx' transaction –initial handling by ST level

	ST level 1						ST level 2						3							
Status	Wait	Wait - Customer	Wait - Implementation	Wait - User	Wait - Vendor	ST level 1 total	Wait	Wait - Customer	Wait - Implementation	Wait - User	Wait - Vendor	ST level 2 total	Wait	Wait - Customer	Wait - Implementation	Wait - User	Wait - Vendor	ST level 3 total	Grand Total
Accepted	990	43	251	3119	247	4650	615	47	197	873	57	1789	140	11	45	225	9	430	6869

Table 4: Wait Status Usage across Support Team Levels

4650 of the 6889 (68%) of transaction logs with 'Wait-xx' sub-status was created by a level 1 ST. Given our overall experience with help-desk/triage processes, we find this metric an anomaly as incidents (SR's at level 1 support level) are typically not expected to need information from customer, vendor or a fix with such a high frequency. Instead, we expected to find more 'wait-xx' sub-status when the request is at level 2 or level 3 support. This raises the question of whether this could be a result of a push-to-front approach where level 2 and level 3 teams may be assigned incidents back to level 1 because of organization expectations? There is also the possibility of this being the outcome of level 1 working under the expectation of having to resolve themselves and not escalating soon enough to level 2 or level 3. We explore this further in our analysis of 'Incident only' dataset.

- Determining if use of 'wait-xx' pattern shows increasing case aging (while still maintaining agreed SLA's) might provide skewed results when assessed at an aggregate level. We decided to undertake this analysis at 'Incident only ' and 'Problem only' datasets.

- Top 15 teams that use 'wait-xx' sub-status most are listed below (~53% of all transactions with 'wait-xx' Status). Interestingly, and re-confirming our interpretation of 'wait-xx' sub-status usage, it is dominated by first level support teams within all teams.

Teams using Wait Status most						
ST	Wait	Wait - Customer	Wait - Implementation	Wait - User	Wait - Vendor	Grand Total
G97	282	2	67	704	1	1056
G96	108		28	226	2	364
G230 2nd	122		18	190		330
S42	49		2	178	43	272
G92	13		10	191	30	244
D5	22		6	212	2	242
D8	9		5	194		208
D2	33		5	92	24	154
S56	18		4	123	1	146
S49	50		3	46	43	142
S43				136		136
D3	54		1	57	13	125
D7	25			97	3	125
D1	57	3		51		111

Table 5: Teams that Leverage 'Wait-xx' sub-status the most

- Ping Pong Behavior Analysis:

To explore ping pong behavior, we used the following a Disco led visualization for the following datasets

i. Master dataset with activity defined by ST level transition and ST transition
ii. Wait user dataset[6] (3501 records) with activity transition same as above

Fig 6 below provides a quick view of org lines that the 9395 SR's were first assigned to. An overwhelming 98% of the cases were handled by Org lines C & A2 (9285 cases in total). Of these 6477 SR's were routed first to Org line C (handled 6888 cases in total) and 1282 cases were routed first to Org line A2 (handling 2397 cases in total). A2 also receives 905 SR's redirected to it from Org line C and in turn passes back 489 SR's. There is limited SR re-routing between other org lines.

[6] This dataset is also available for reference in Appendix 1

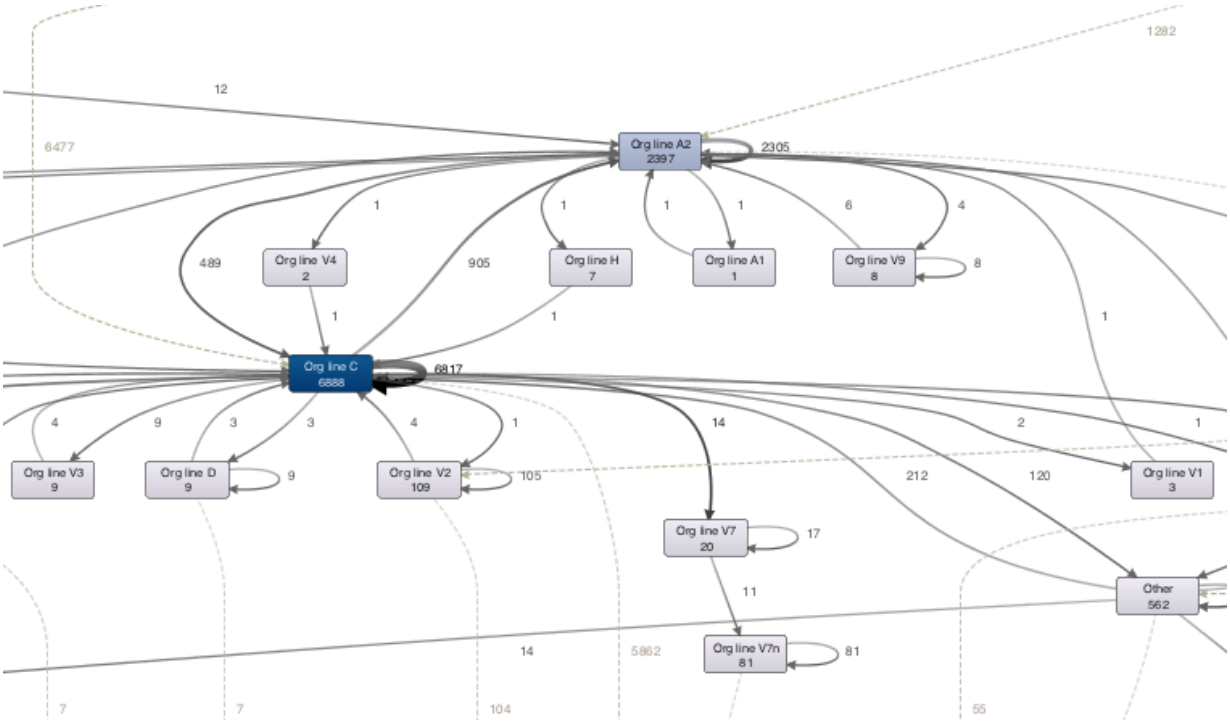

Fig 6: SR initial assignment at Org line level

- To determine if there is a Ping-Pong pattern at the Org line level, we tried to deep dive and determine if cases routed from Org lines C to A2 being re-routed back and vice versa? First, we filtered the 905 SR's routed from Org line C to A2 and visually analyzed this

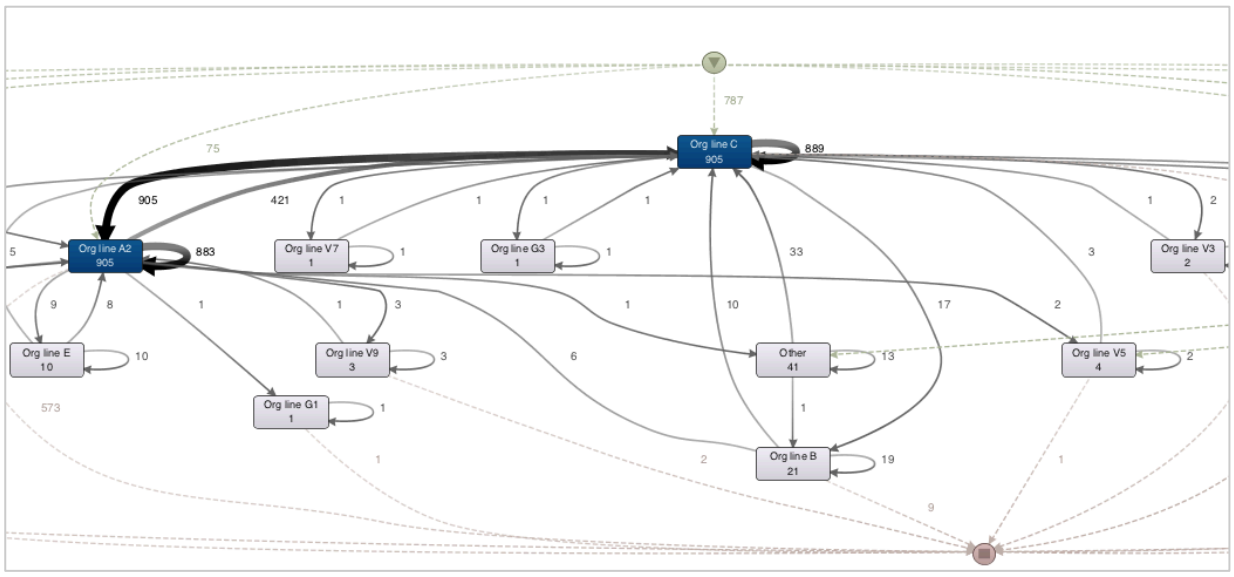

Fig 7: flow analysis of 905 SR's between Org line C and A2

- Org line C received 1280 SR's (inclusive of duplicate SR's) in total, 787 directly routed and 493 routed from other org lines. Since it only handled 905 discrete cases. It implies that there are a total of 365 common (or 'Ping-Pong') SR's between Org line C and other Orgs. There are only 16 potential Ping-Pong cases with other org lines (stats in table 6 below), suggesting almost 349 Ping-Pong cases between Org line C and A2.

Org Line C				
Given to	905		Received from	'ping pong' potential?
G3	1	1	G3	1
V7	1	1	V7	1
A2	905	421	A2	
V11	0	15	V11	
E	0	5	E	
V3	2	1	V3	1
V7n	3	3	V7n	3
V5	0	3	V5	
Direct	0	787	Direct	
B	17	10	B	10
Other	0	33	Other	
	929	1280		

Table 6: All SR's handled by Org line C (Incl. potential ping pong SR's)

- Further, we explored the SR transitions at functional unit level. Org line C consisting of C_1 to C_7, E_1 to E_10 and V1_V3. Org Line A2 comprising of A2_1 to A2_5, D_1 to D_3. When we explored the dataset, we noticed approximately 3357 transactions having Org line C had involved ST functional Div. as A2_1. This represents 0.06% of the transaction. We assumed that this would not impact our analysis at the functional team level. There were also 9534 transactions with blank values in Functional division. Interestingly (as can be seen in fig 8 below) this particular 'blank' function receives and routes a significant number of SR with A2_1 (288 & 267).

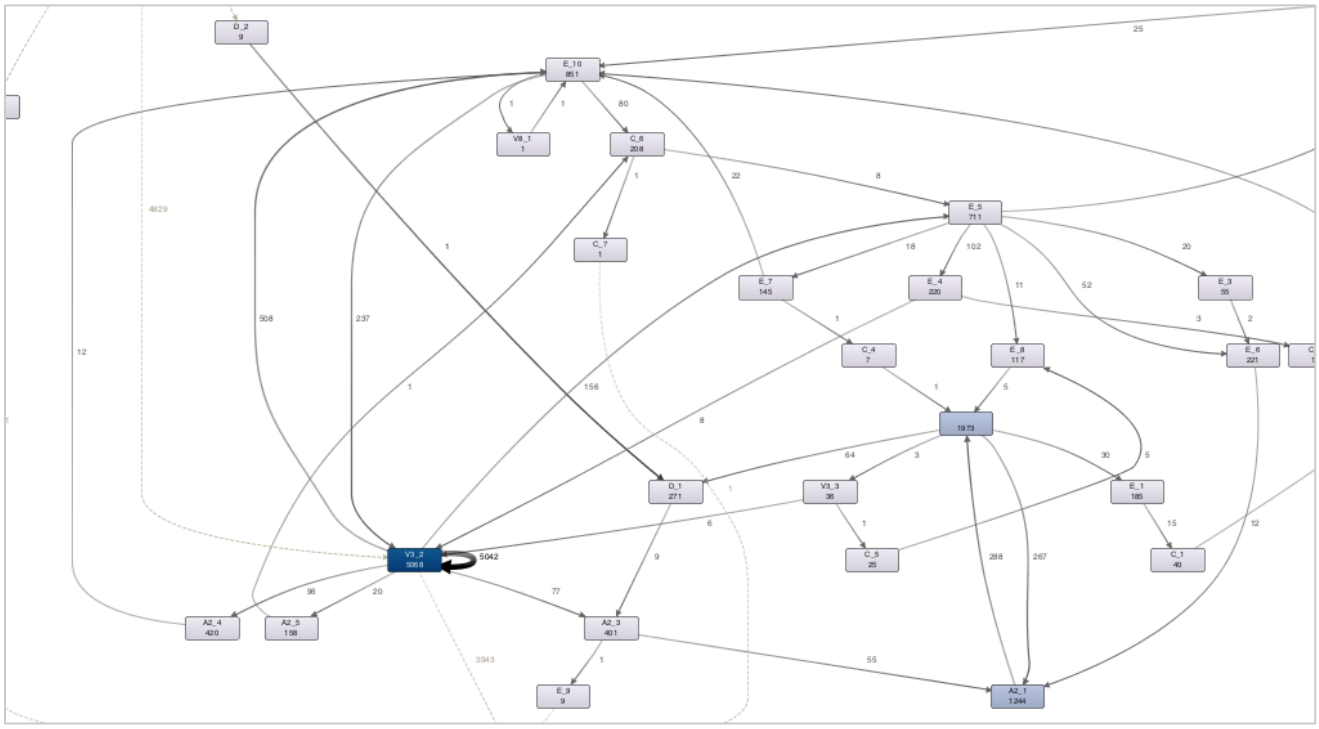

Fig 8: SR flow at functional div level (partial view only) to highlight potential 'Ping-Pong' cases.

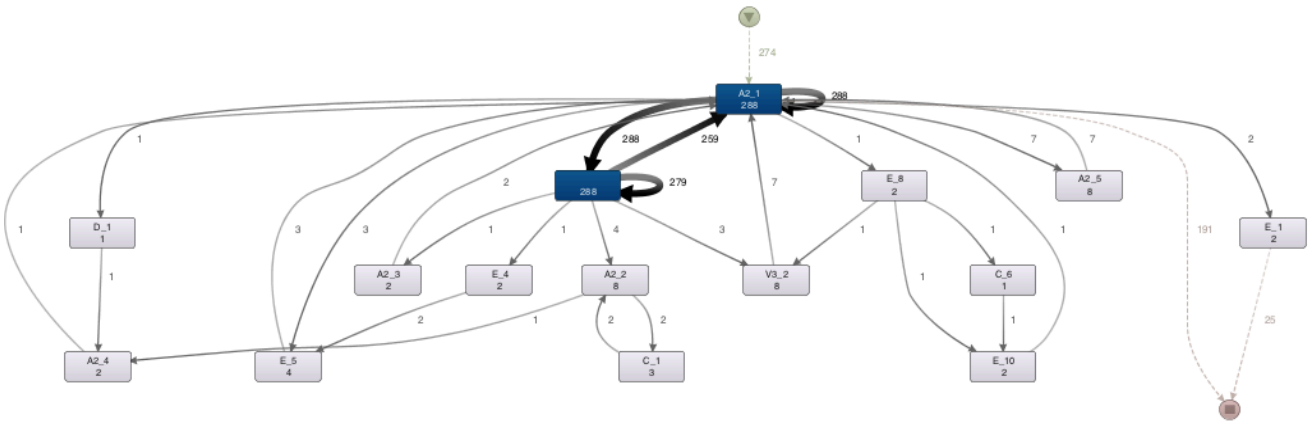

Fig 9: Filtered view of potential Ping-Pong cases to check if additional flow patterns emerge

A quick check at the dataset revealed 295 SR's routed to both 'Blank' and A2_1 Functional division [7]

On further examination, multiple org lines are associated with this 'blank' functional division with the maximum SR's being handled by V7n (75) and V11 (112). In analyzing a matrix of SR's handled by A2_1

[7] Case-list provided in Appendix 1

and 'Blank' functional division, we find A2_1 being associated with multiple org lines (A2, B& C). This complicates our ability to recommend concrete actions, but is a large enough set to explore data quality improvement.

- At the ST level, we found potential ping –pong patterns between D4 and N26 2nd (~51 SR's cases) , D8 and G179 (~40 SR's) , and S42 and S43(~ 16 SR's). As a general observation, we see low pattern of cases transitioning between Support teams at the same level vs. cases transitioning between teams at different levels.

Analysis of Incidents only dataset

- The goal of any incident management process is quick, satisfactory resolution at first contact, generally measured by an FCR metric. Taking our incident-only dataset and passing it through Disco, we undertook some quick visual analysis.

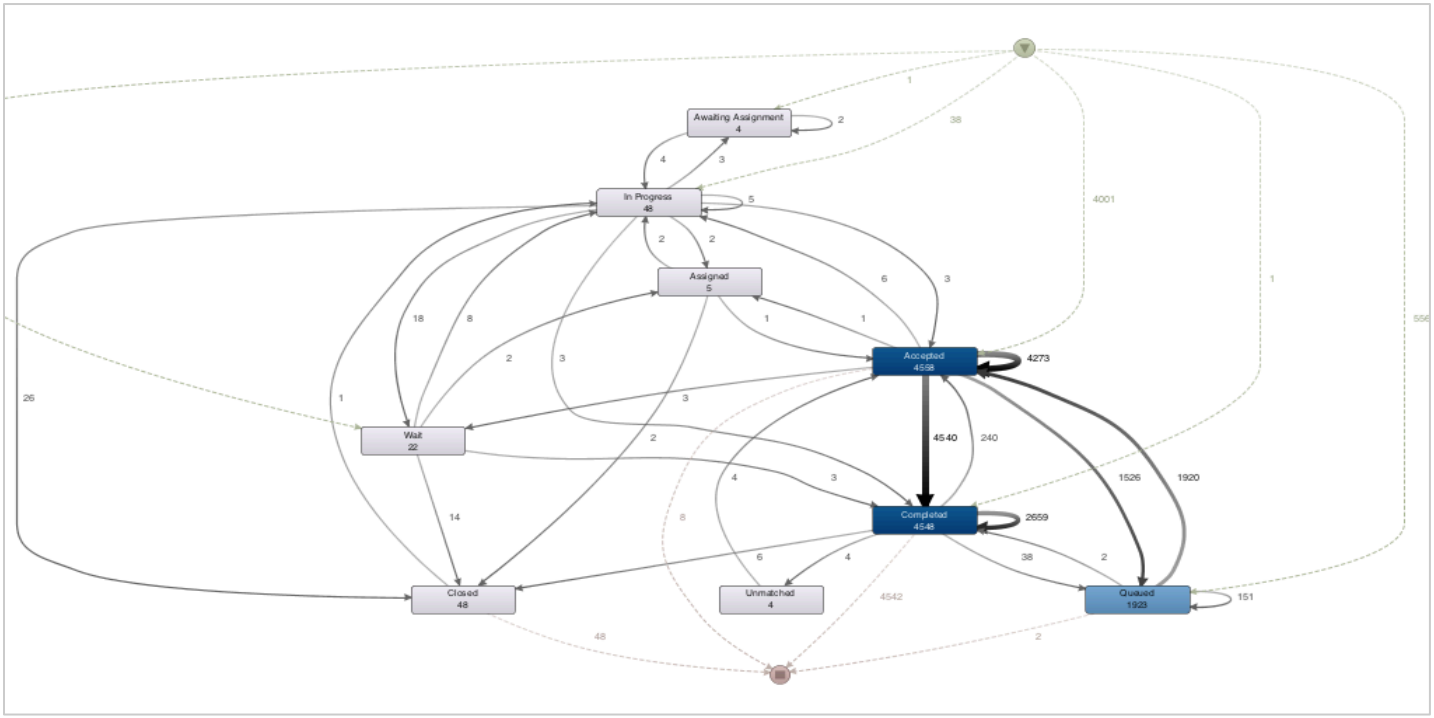

Fig 10: Flow analysis of incidents with activity as 'Status'

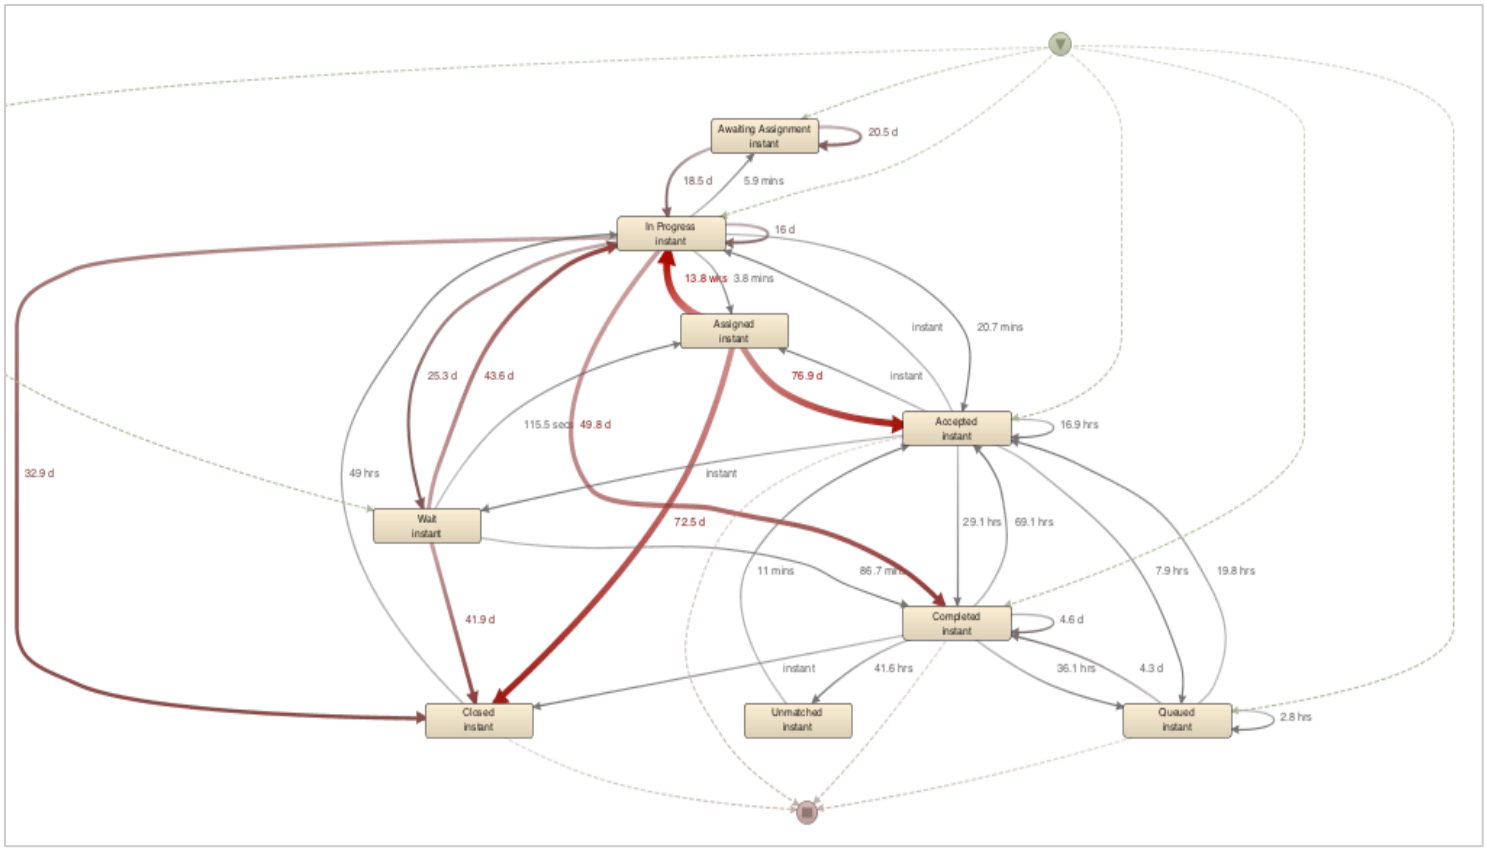

Fig 10: Performance metrics, Incidents only

	High	Low	Major	Medium
SR Count	51	2254	5	2290

Table 7: distribution of Incident only SR's across Impact levels

- Key observations:
  - Distribution pattern in table 7 above reflects our expectation. Hence, the usage of 'wait-xx' sub-status is further paradoxical unless a whole set of user unique incidents are being logged.
  - Analyzing at 'status' level, 4001 SR's had 'accepted' as an initial transaction status, 566 SR's had 'queued/ awaiting assignment' as initial transaction status.
  - Of 4606 incident only cases, 4548 were completed at an average time to resolution of 69.1 hrs. (~ 3 business days). This looks like a potential area of improvement, as this average resolution time is probably too long for low and medium impact cases.
  - Only about 48 cases are closed, showcasing perhaps a desire for problem /RCA for the remainder of the completed cases. This could also reflect an ambiguity in using the right status levels and might be a training issue.

- We used status/sub-status combination to dive deeper into the completed case stats. Here is what we discovered:
    - About 1954 cases progressed to closure rapidly with status as 'Complete/in-call'. Of this 104 SR's regressed back to 'Accepted/In Progress' status. Nonetheless, this data showcases a ~40% FCR metric. The goal for the VINST IT team should be to raise this metric. Average time to resolution of these cases was 37.9 minutes, which reflects a potential opportunity to leverage this channel.
- Approximately 1796 SR's moved to Completed/Resolved within 8.6 Hrs. This seems in line with expectations although it is difficult to judge without the 'urgency' values in the dataset.
- Approximately 1530 SR's used an interim 'wait-xx' sub-status between accepted/ in progress and completed/resolved. A majority of these (~1098 SR's) used 'wait-user'. We decided to deep dive on these SR transactions. We found significant back and forth between multiple 'wait-xx ' sub-statuses (please see fig 11 below). Our recommendation is for the VINST team to analysis these 1530 SR's closely and determine approaches to improve first call resolution.

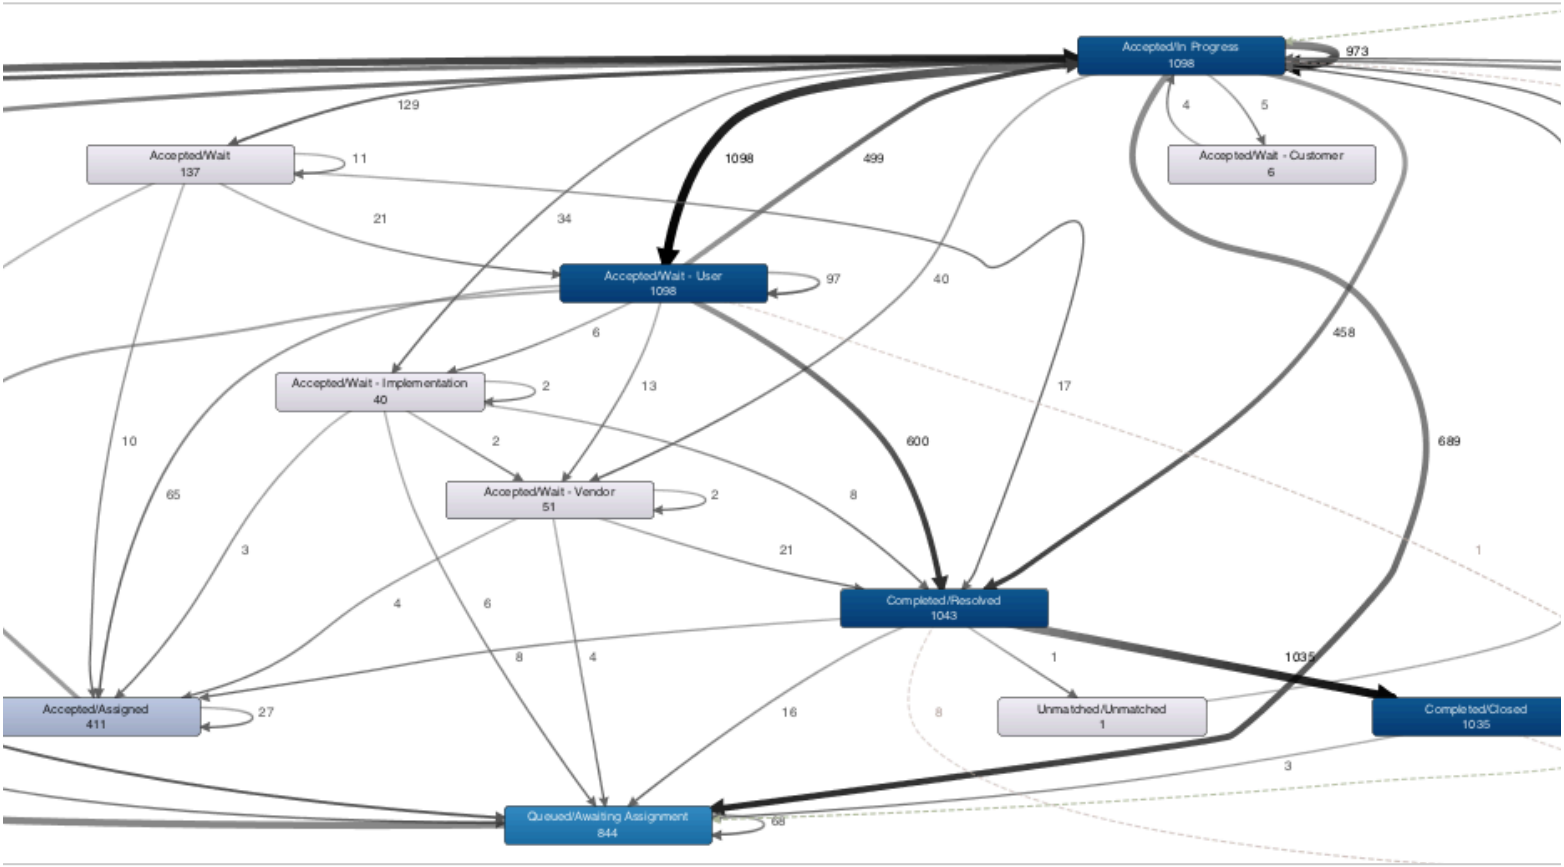

Fig 11: transaction flow for 'Accepted/In Progress ' cases

Analysis of Incidents to Problems dataset

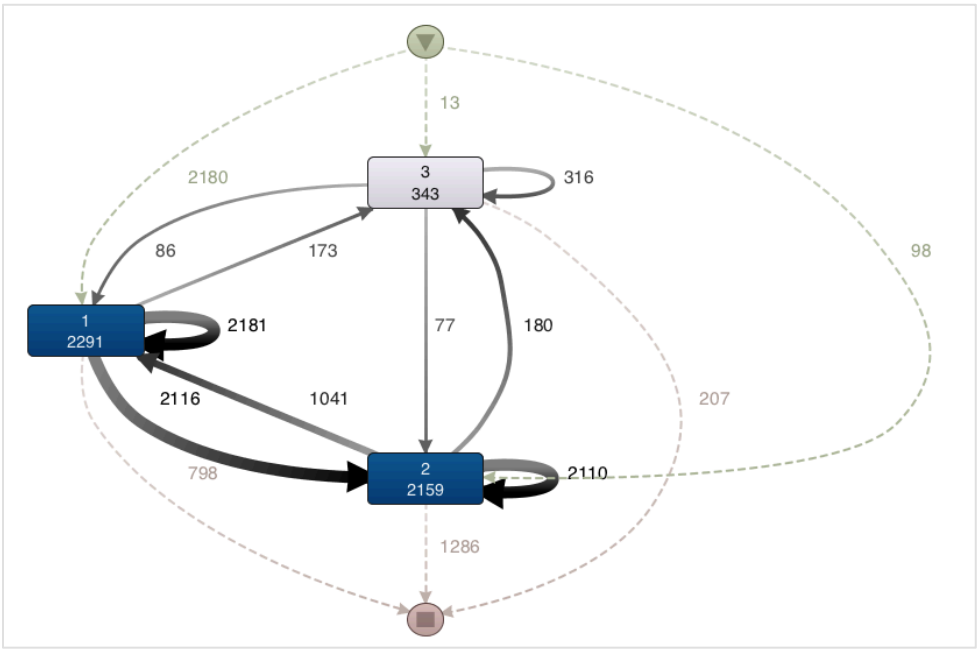

Fig 11: ST level flow pattern for Incidents to Problems SR's

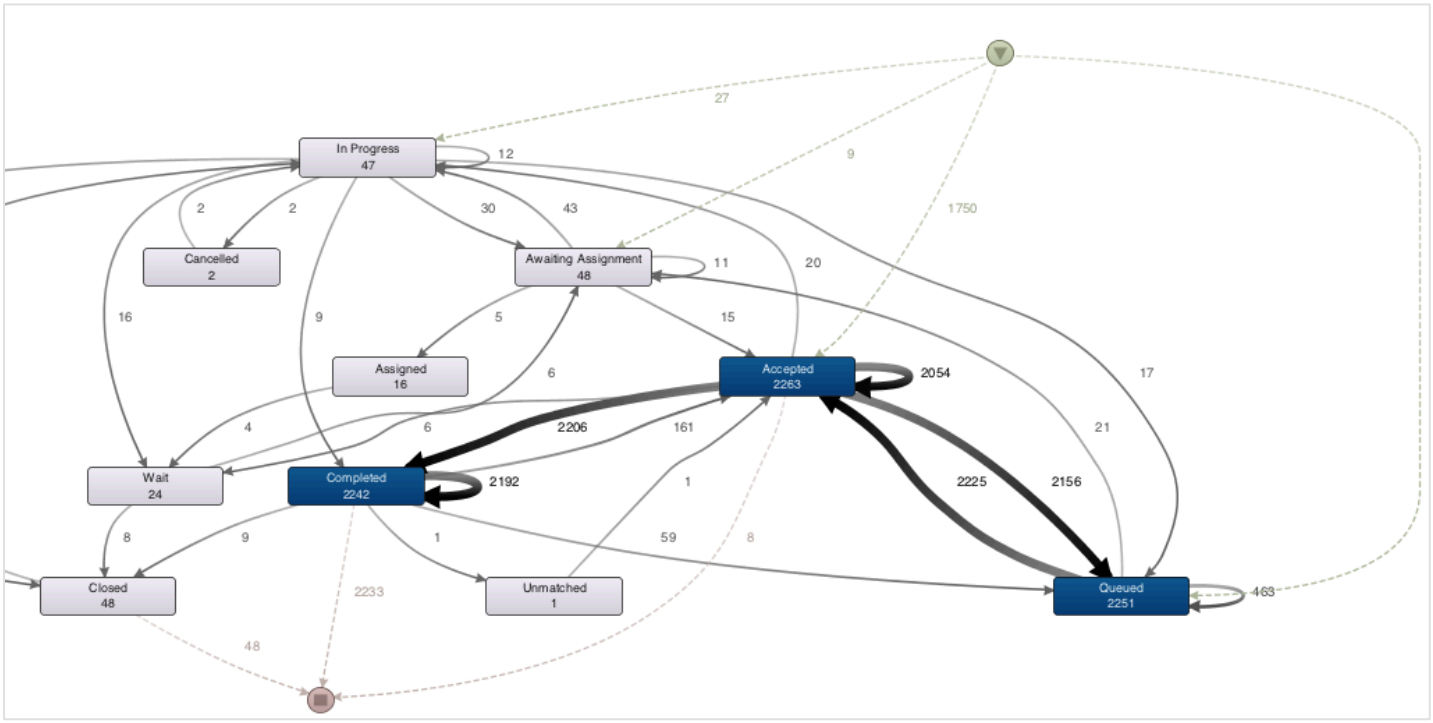

Fig 12: SR flow (case frequency) by status

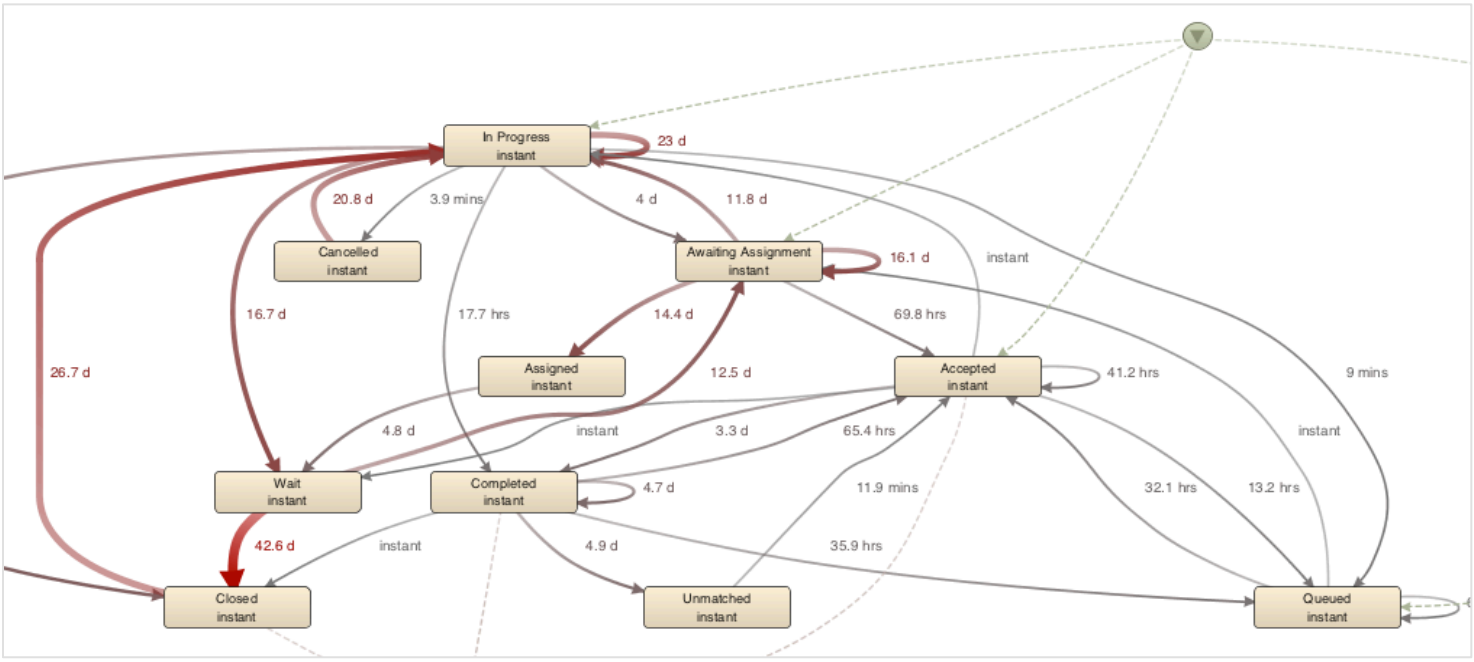

Fig 12: SR performance metrics

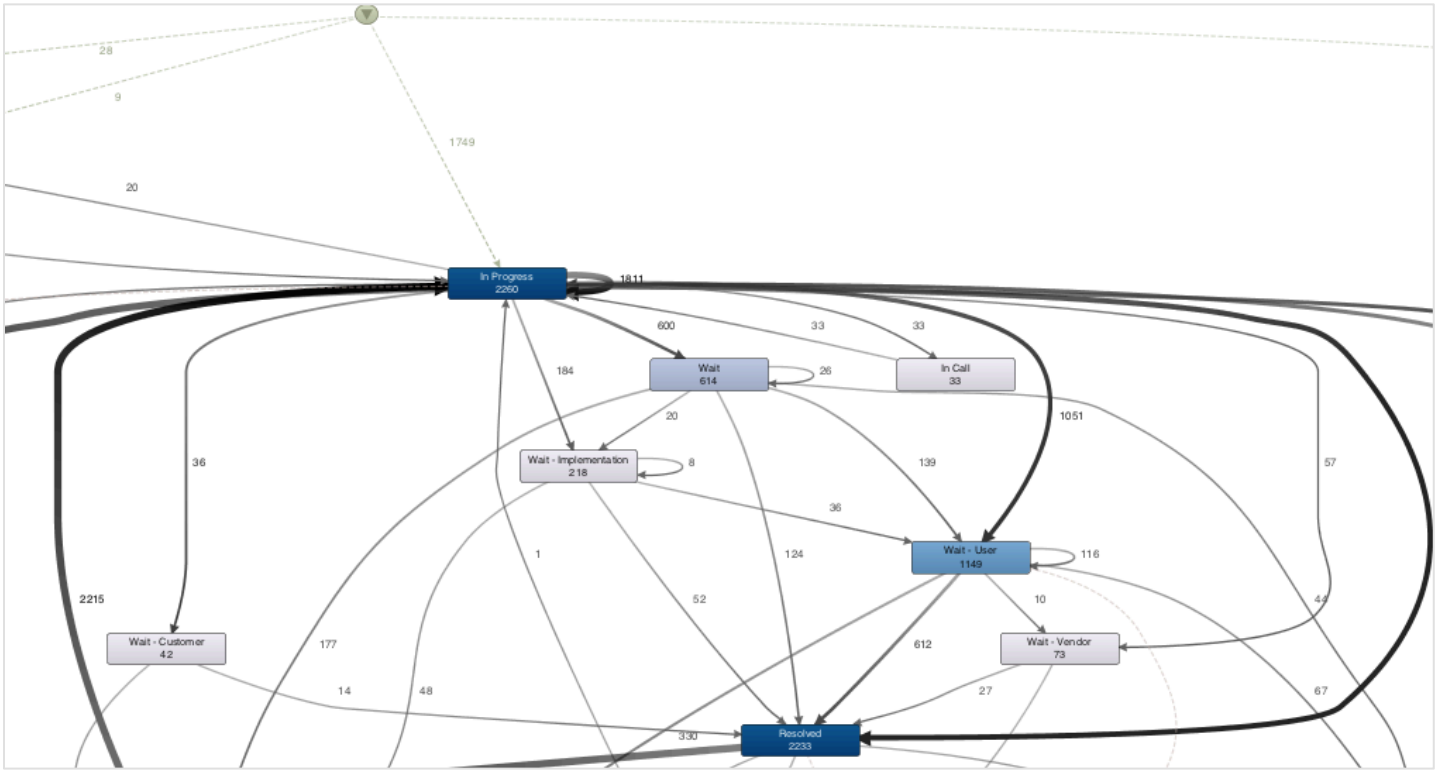

Fig 13: 'Wait-xx' usage pattern

	High	Low	Major	Medium
SR Count	148	802	5	1335

Table 8: Incidents to Problems dataset distribution by impact

- Observations:

  - 2290 / 9395 SR's constitute this dataset ~ 25% of the overall SR pool. We recommend limiting this by improved training of first level teams, given the high incidence of low and medium impact SR's in this bucket.
  - There were insignificant major impact SR's assigned escalated from level1 teams to 2 or 3. Much of the escalated SR's had medium impact
  - Average time to resolution once a case has been accepted in 3.3 days. The typical case flow pattern is accepted –>Queued->Accepted ->completed. This is understandable as the cases transition from one ST level to another. Overlaying state transition metric with case aging metric will provide an enhanced view of bottlenecking stages.

Analysis of Problems only dataset

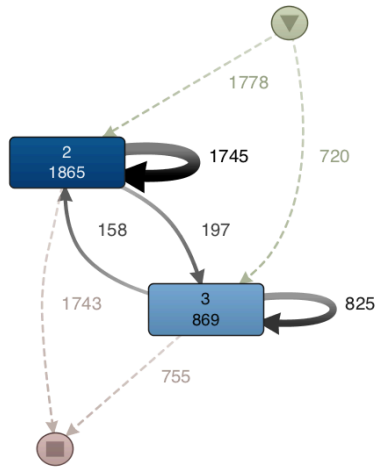

Fig 14: ST level flow pattern for problem only SR's

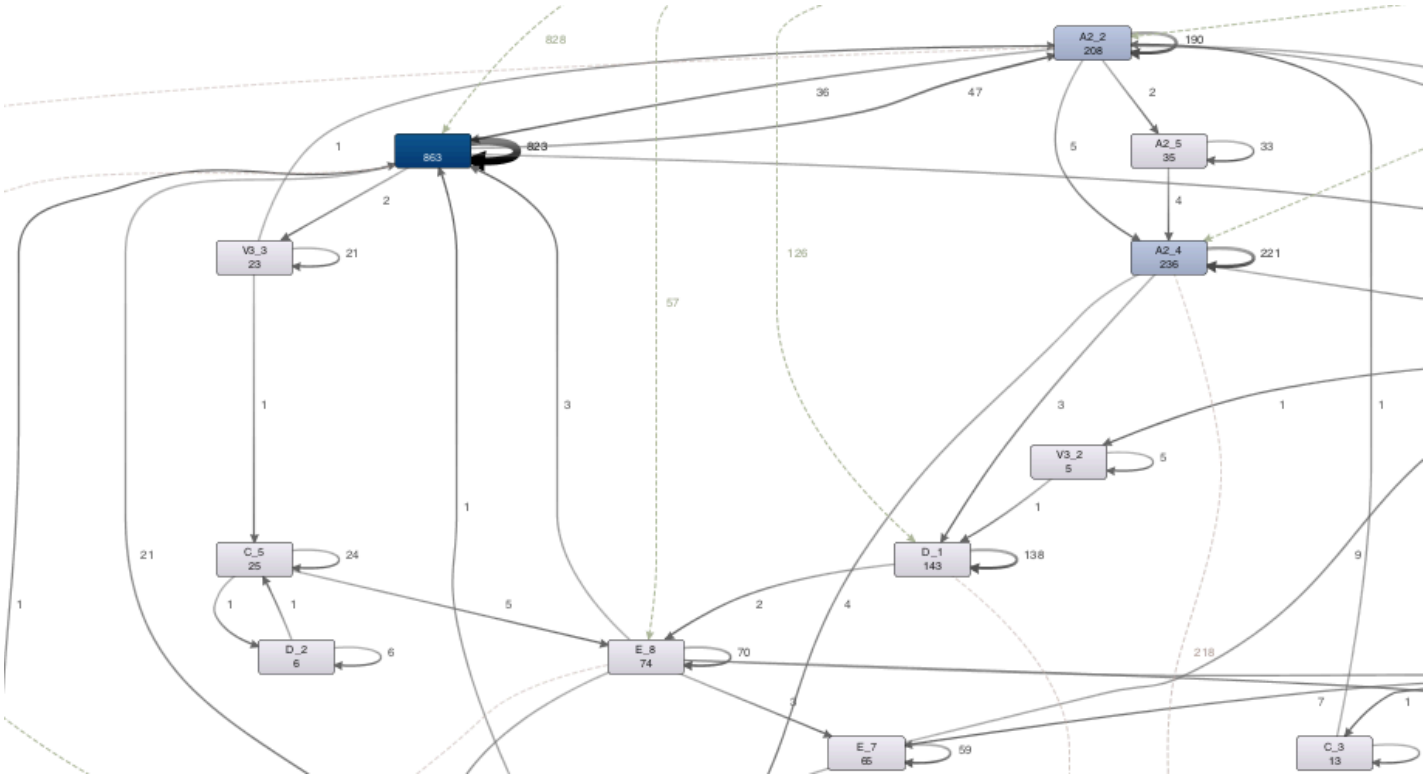

Fig 15: ping –pong analysis between functional divisions.

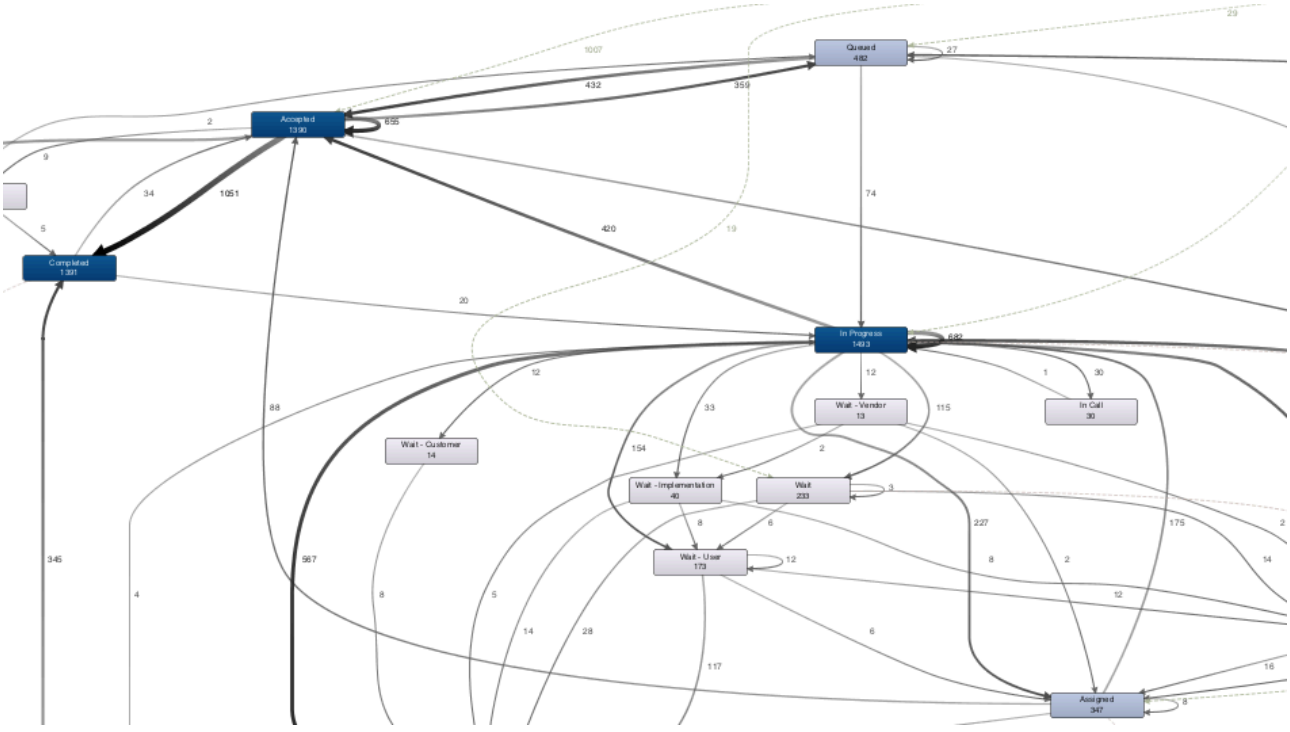

Fig 16 – ‘Wait –xx’ status usage

Observations:

- This dataset comprised the highest incidence of 'Major' impact cases (135 out of 2498 SR's). This is as expected.
- Routing of SR's to the right level is also much cleaner, with lower relative incidence of 'Ping-Pong' behavior between levels.
- We do not observe high incidence of 'Ping-Pong' behavior between functional divisions. This is most likely because initial assignments / routing is clear.
- 'Wait –xx' usage. Relative to Incidents only, we observe a much lower incidence of 'Wait-xx' status usage. Only 226 out of 2498 SR's go through this stage. Also, this is used post 'in-progress' sub-status transition showing empirically a more involved decision in using this sub-status.

Summation and key findings

Across the 4 datasets analyzed here are our observations of a few consistent patterns. We have limited our assessments at process levels and felt a need for a next level of statistical analysis, which we unfortunately could not undertake. We have also not analyzed in detail at the product, owner and ST level with the overall philosophy that behavior patterns at those levels reflect reactive symptoms and will correct themselves if overall process patterns are improved. It is also our philosophy that the goal of process analysis should be to detect areas for process improvement vs. identifying low-level behavior patterns. Based on our observations, we recommend the following three steps:

1. Consider adding a triage process upfront that helps distinguish an incident from a problem.
2. Consider different state transitions for incidents and problems. Also establish a cleaner pattern between status and sub-statuses. This might involve rationalizing the current status/ sub –statuses.
3. Consider improved training and handling of Level 1 support teams with a goal of improved FCR.

References

- Discovery, Conformance and Enhancement of Business Processes: van der Aalst, Wil M. P.