



**VIRTUAL EXPERIENCE
OCTOBER 11-14**



Up Your Uptime With Automation

A Technical Paper prepared for SCTE by

Nancy McGuire

Executive Director – Reliability Engineering, Operational Support
Comcast Cable
1800 Arch street, Philadelphia, PA 10103
(215) 286-2290
Nancy_McGuire@cable.comcast.com

Kathy Fox

Vice President – Product Management, XOC
Comcast Cable
609 Odin Rd, Coudersport, PA 16915
(856) 912-7850
Kathy_Fox@cable.comcast.com



Table of Contents

Title	Page Number
1. Introduction.....	3
2. Getting Started – Process Alignment.....	3
3. The Results of Process Alignment.....	5
4. Additional “Sameness” Opportunities.....	7
5. Align the Process, Then Automate – One Way	8
6. How Covid Impacted Automation Activities	9
7. Necessity: The Mother of... <i>Automation!</i>	9
8. How We Automated It	9
9. Conclusion.....	16
Abbreviations	17

List of Figures

Title	Page Number
Figure 1- Traditional Division “Node Soak” Process.....	4
Figure 2- Division 2 Dashboard Tasks canceled based on triage	5
Figure 3 – Division 3 Dashboard Tasks Canceled based on triage.....	6
Figure 4 – Breakdown of One Division’s Plant Maintenance Cancelations – Daytime hours	6
Figure 5 – Breakdown of Another Division’s Plant Maintenance Cancelations – Daytime hours.....	7
Figure 6 – Multi-Soaking Outage (MSO) Repair Results.....	8
Figure 7- volume reduction via correlation.....	11
Figure 8- Pre-portal Workflow	13
Figure 9- Post-portal Workflow.....	14
Figure 10 - Mean Time Graphs.....	15

1. Introduction

This is being written in the summer of 2021, a year-plus since the global pandemic that sent many of us home to work, and created an entire category of 2020 memories – “my/our COVID project” – ranging from sourdough bread to bathroom remodels to fitness journeys. If there was such a thing as an unofficial, job-related COVID project, automation would assuredly qualify! From plant maintenance to residential product support, circumstances arose over the past year that catalyzed us to align processes and achieve automation.

Maybe this feels familiar to you: Volumes upon volumes of tickets. Late night/early morning pages and dispatches. On-call team member burnout. All are symptoms of large, complex systems, built and patched and rebuilt over decades, and probably showing multiple signs of wear. You, of course, “just” want to ensure you know what is going on, where, and when.

But how do you sort what is truly an impacting event that requires immediate attention, from the what can be a daunting level of “informational noise”? How do you get the right eyes on the right problem at the right time?

The answer that worked for us was a duet of process alignment and automation. In this paper, we will review some of the automation opportunities we identified, how we identified them, and the impact they had on our teams and business operations functions. We will review two different approaches to address these opportunities – informational noise reduction, and optimization by correlation -- in two quite different areas of our operation: first, by our field plant maintenance technicians, supported by our XOCs (eXcellence in Operation Centers), and second, in our national Residential Product Support team.

2. Getting Started – Process Alignment

In consolidating and post-consolidating environments, a necessary course of work is to find “sameness,” for purposes of scale. For our national and field teams, we have been focusing on sameness across all functions. However, overall “sameness” only delivers that initial goal of industry consolidation and geographic clustering -- scale. As important as scale, and the next step after achieving it, is internal synchronicity. We define that here as “the best of best practices, applied at scale,” across the consolidated organization. Internal synchronicity optimizes how things are done, and it does so at scale. It’s why we’re on a quest, partly characterized in this paper, to identify best practices/processes/ solutions and implement them everywhere.

As we identified areas of differences, an initial area of activity we believed could serve large benefits to customers and employees was identifying and implementing best practices around node outage treatments. When we dispatch plant maintenance technicians for node outages during the day, we are pulling them off other work that is in their queue. When we dispatch plant maintenance technicians after hours, we incur overtime, and when there is enough on-call activity, those technicians can burn out. It is therefore important that we dispatch them only to nodes where they can take action to fix problems. It is about getting the right technician to the right place at the right time. It is worth noting, plant maintenance technicians are the more expensive technicians, with the pricier bucket trucks, as compared to fulfillment technicians. And we all know that when dispatched at night, technicians deserve for it to be a legitimate issue.

We needed to determine what steps we could take to reduce the number of times plant maintenance trucks get dispatched, only to find that the node was out because the power was out, or some other event the technician was not able to repair – the classic “NTF” (no trouble found). We have three different divisions and as we looked at node outage practices, we found different processes for each of the divisions.

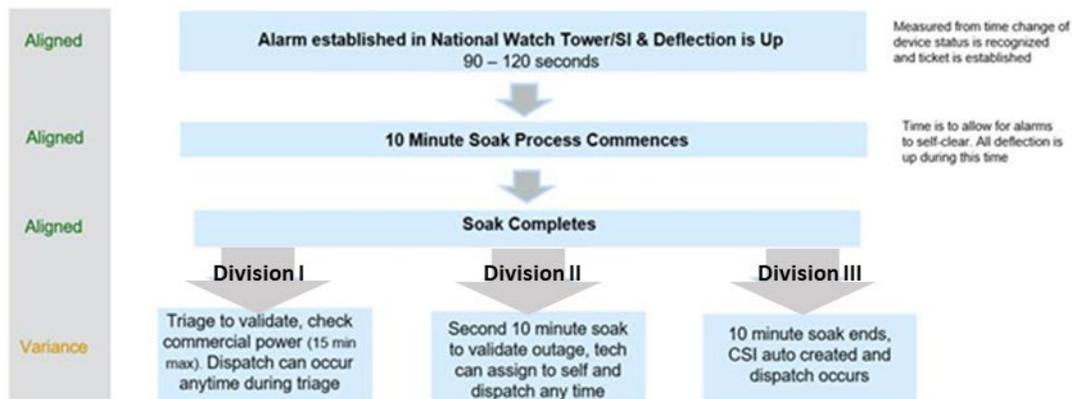


Figure 1- Traditional Division “Node Soak” Process

As we see in Figure 1, above, there was alignment in three areas, but variance existed once a node completed the initial 10-minute “soak,” a term we use to indicate the period of time during which node components self-clear any alarms. It is important to note that nodes go into soak when a certain percentage of modems served by that node go through a registration state change (to offline).

The division teams came together to review their processes and resulting data, comparing apples to apples, to see who had the most efficient plant maintenance dispatch records, based on total dispatches, fix codes, and NTFs. Perhaps not surprisingly, folks can have a good bit of passion for their individual tools and/or processes. So, after a notable amount of review of data and discussion of the data, the node outage process utilized by Division I was identified as the best solution for the customer, the employees, and the company. We referred to the activity as “outage pre-verify,” because the XOC was verifying that the outage exists before dispatching staff. Once we had a process name and agreement, the team then focused on developing a thorough implementation plan which included detailed training and communication plans; a review of the LOQs (Lines of Questioning) used by the Division I team during the triage activities; determining and preparing for expected HFC (Hybrid Fiber-Coax) desk task increases because of the triage activity; and expected truck roll decreases related to plant maintenance. Once we had developed this detailed plan, we launched it on a region-by-region basis across divisions 2 and 3. In some areas, we launched more quickly than originally planned, because it was such a good story, from a plant maintenance perspective.

3. The Results of Process Alignment

How has the Outage Pre-Verify work been received by the two other division field teams? Charles Detwiler, a Comcast Maintenance Supervisor in the Pittsburgh area summed it up this way:

“Outage Pre-Verify has reduced overtime and man hours drastically. HFC techs are putting information in tickets, so we now have a better understanding of what we are going to be troubleshooting and there is also much more trust that this is a legitimate outage. We just had a regional call and reviewed the significant truck roll reductions since rolling out Outage Pre-Verify. Lastly, our techs are less stressed because we are rolling fewer trucks at night, so we are getting more rest.”

Steven Musembi, VP (Vice President) of XOC for Comcast’s Central Division credits the introduction of Outage Pre-verify as *“a direct contributor to node performance in the Central division. OPV (Outage Pre-Verify) freed up labor via transaction reduction to help find and address plant issues, improving the customer experience for all of our customers.”*

Issues that are detected via our tools and rule sets, which become tasks for our HFC technicians, are referred to as Dashboard Tasks. Whenever there are node issues that hit our thresholds, that creates a dashboard task. The numbers speak volumes. Figures 2 and 3, below, show the total number of dashboard tasks and the number that were cancelled as a result of XOC outage pre-verify work. The overall result was reduction of all dashboard tasks by nearly a third (29.32% and 28.5%.) Internally, we were also able to break these out by daytime and after-hours Plant Maintenance dispatches, to show the overtime savings. Again, these reductions happened and continue to happen because the HFC desk technicians are following through on the same set of OPV (Outage Pre-Verify) LOQs (Lines of Questioning) to properly triage a node outage.

OPV Details - Division 2 Fiscal June 2021			
Region	Total Dashboard Tasks	# Cancelled Dashboard Tasks	% Canceled Tasks
1	14765	3847	26.05%
2	11679	3620	31.00%
3	8105	1624	20.04%
4	7807	3326	42.60%
Totals	42356	12417	29.32%

Figure 2- Division 2 Dashboard Tasks canceled based on triage

OPV Details - Division 3 Fiscal June 2021			
Region	Total Dashboard Tasks	# Canceled Dashboard Tasks	% Canceled Tasks
1	7487	1974	26.30%
2	5582	1962	35.15%
3	3843	978	25.45%
4	4241	1193	28.13%
5	2134	532	24.93%
Totals	23287	6639	28.51%

Figure 3 – Division 3 Dashboard Tasks Canceled based on triage

Daytime Savings – Division II

Region	# of Dash Tasks between 8 am and 5 pm	# Canceled WT Tasks between 8 am and 5 pm	% of Canceled WT Tasks Relative to the # of Dash Tasks (8 am to 5 pm)	Total Tme Savings in Hours (Tasks 8am-5pm *1.5 hours)
1	5,649	1,451	25.69%	2,176.5
2	4,610	1,357	29.44%	2,035.5
3	3,850	820	21.30%	1,230
4	3,387	1,408	41.57%	2,112
Total	17,496	5,036	28.78%	7,554

Figure 4 – Breakdown of One Division’s Plant Maintenance Cancelations – Daytime hours

Daytime Savings – Division III

Region	# of Dash Tasks between 8 am and 5 pm	# Canceled WT Tasks between 8 am and 5 pm	% of Canceled WT Tasks Relative to the # of Dash Tasks (8 am to 5 pm)	Total Tme Savings in Hours (Tasks 8am-5pm *1.5 hours)
1	2,924	852	29.14%	1,278
2	2,343	888	37.90%	1,332
3	1,672	507	30.32%	760.5
4	1,661	563	33.90%	844.5
5	911	239	26.23%	358.5
Total	9,511	3,049	32.06%	4,573.5

Figure 5 – Breakdown of Another Division’s Plant Maintenance Cancellations – Daytime hours

4. Additional “Sameness” Opportunities

The Division I team also shared its operating practice for MSOs (Multi-Soaking Outages). A multi-soaking outage is the term used to describe a node that drops into soak, which as a reminder occurs when a certain percentage of modems associated with that node experience a registration state change, indicating that “something is up” at the node. What’s different in a multi-soak outage is that the node self-clears before the 10-minute soak time elapses, but alarms and re-soaks itself again, in short order. The division’s process was to triage and dispatch help to nodes that dropped into soak, and then cleared, four or more times in any rolling 24-hour period. Because we were going to be reducing plant maintenance truck rolls with our Outage Pre-Verify work, it made sense to also implement this same practice across all three divisions. We knew that implementing the MSO functionality would drive additional HFC technician tasks as well as additional plant maintenance-related truck rolls. However, the HFC leadership teams were gaining efficiencies by reducing tasks in other areas, and the Plant Maintenance teams were very capable of taking on this work, given the significant reductions in truck rolls as a direct result of the Outage Pre-verify work.

Figure 6 shows that our plant maintenance teams are, on average, finding and correcting node multi-soak issues more than half (57.5%) of the time. Again, these are intermittent outages that we are correcting before they become a larger outage – a big lift for the customer experience.

Division II	Region	Tasks	% Region
	1	137	54.80%
	2	102	72.30%
	3	160	57.60%
	4	77	72.00%
Division Total		476	61.30%
Division III	Region		
	1	177	48.10%
	2	136	58.10%
	3	145	63.60%
	4	132	54.30%
Division Total		590	57.50%

Figure 6 – Multi-Soaking Outage (MSO) Repair Results

With the solid progress being made for MSO with the setting of 4 times in a rolling 24 hours, we are poised for the divisions to move to 3 times within 24 hours, to drive even more reliability into the infrastructure – which automatically means fewer outages for our customers! One division has already moved to 3 times in 24 hours, while our two other divisions preparing to make the move as well.

5. Align the Process, Then Automate – One Way

Once the teams aligned on the process, we needed to lighten the load for our HFC technicians, who had taken on a great deal of additional triage activities. And now that we were all on the same process, we could build consistent automation. All those canceled jobs were canceled because of the excellent work performed by these HFC technicians – and we did not add headcount to our HFC desks to support the outage pre-verify or the MSO work. Together, the division and national team developed a user story to capture a description of a software feature from an end-user perspective. This allowed us to reduce “eyes on glass” time for the HFC desk technicians. Considerations for the user story included:

- Utilizing automation and other tool inputs to validate node outage events and reduce pre-verify questions, using already-available data for:
 - Verification of device registration state
 - Confirmation of no upstream activity – flatline node
 - Validation of power supply status
 - Determining “storm mode” status
 - Verification of headend or CMTS (Cable Modem Termination Server) alarm
- Saving an estimated 3 minutes per pre-verify task (on average, a pre-verify task takes 7 minutes prior to the automation), by reducing the number of manual actions needed for task completion
- Saving an estimated 30K+ hours per year at the XOCs, on eliminated workflows.

This automation work is in progress for implementation in 2021. Future automation enhancements will include checks of Xfinity Home devices served by a node, to determine connectivity type (backup / cellular), to further the accuracy of automated checks related to commercial power availability.

6. How Covid Impacted Automation Activities

When you are supporting the flagship products for a major cable/broadband company, it is imperative that they work at or better than the expected 99.999% of the time. This certainly made 2020 a year to remember! The pandemic required teams and companies to shift perspectives, cultures and make some really challenging decisions. Employees were having difficulty with the lack of in-person social engagement, hallway conversations and the biggest issue, the blurring boundary between work/life hours.

7. Necessity: The Mother of...Automation!

We had to find those solutions fast! But first we had to identify what could be automated. We looked at redundant tasks such as pulling logs from devices, ticket entry, ticket assignment and triage. Previously, our offshore team did the ticket creation, while another offshore team did the initial triage. Outsourced resources also played a major part in “air traffic control” - assigning tickets to the correct teams for mitigation, as well as escalating to on-call resources when needed. Fortunately, our teams were slightly ahead of the curve here and had been doing the data analysis prior to this occurring. This event just fast-tracked the groundwork we had already laid out.

Having a head start in finding and implementing automation drastically reduced our fears about team burnout and not being able to support our products at the same level. While there were obvious candidates, we were still faced with the challenge of *how* we find automation opportunities.

The undisputed answer: data analysis and best practices.

Specifically, we built reports to understand what our biggest drivers were, and what could we do to streamline our workflows. We also looked for “low hanging fruit,” or, the quick wins. We engaged with our engineering teams to dive into what was important and what was lower priority that could wait until business hours. We worked to create a portal so that our internal customers could create tickets on demand and escalate if needed. We worked on auto-assignment, auto paging and correlation. We developed “soak” periods for alarms and alerts, to give electronics a grace period to reset themselves. Not only did we work to make our processes and systems more efficient and user-friendly, but we also worked to make them smarter.

We also needed to understand who our internal customers were and what needs they had. Any solution we developed needed to be user-friendly and convenient. What existing tools could we share with them to help them do their jobs, and potentially prevent the escalation in the first place? What opportunities are “hiding in plain sight,” for anomaly detection and/or reductions in MTTM (mean time to mitigate)?

8. How We Automated It

In March of 2019, we started taking a hard look at the volume coming into the Production support team. We broke it down by the primary service, looking at each individually. We ran ticket reports to understand volume, trends, and commonalities. Our highest volume related to an email platform. In March of 2019, there were over 74,000 tickets created for this platform alone. Once we did the analysis, we discovered that a majority were auto-resolving, when the main event resolved. In other words, we

were generating tickets for events where they were not the primary issue, but rather an impacted service. It was informational noise.

One of the other findings was that, the outsourced team was largely serving as a pass-through for tickets. An alarm or alert would trigger and be assigned to the Tier 1 queue. The Tier 1 team would respond by reassigning the ticket to a Tier 2 team. Ahem: If no investigation or triage is happening, then why is a ticket entering any queue?

These top-two findings guided our actions and next steps. Our first action was to correlate events with primary drivers. As an example, if one VM (Virtual Machine) went into an alarm state and there were 50 affected applications on that VM (Virtual Machine), we will generate 51 tickets. In the updated noise reduction effort, we now have 1 ticket that has 50 events appended. This reduces volume while still tracking impact. It significantly reduced overall volume and allowed us to find and trend the truly impactful events before they became major incidents.

Based on our analysis of the data, we were able to identify and define this correlation solution, which we presented to our engineering partners. We wrote out exactly what problem we wanted to solve, and what our desired outcome would be. Once we had these defined, we were able to figure out the actions needed to get from the problem statement to the solution. We then presented this to our engineering teams to kick off our efforts and garner buy-in. During this collaboration process, we also developed a deeper understanding of the platforms we support, strengthened our relationship with the engineering partners and most importantly, provided results. This also allowed the engineering teams to focus on the stability of the product rather than sort through a high volume of informational noise.



Figure 7- volume reduction via correlation

The impact of this implementation also had other benefits. By limiting what alerted to the affected service, application, or infrastructure, the mean-time-to-mitigate was reduced and customer NPS (Net Promoter Score) saw a rise in positive feedback.

“I lead an engineering and operations team for enterprise-scale complex applications. Large events would often lead to a storm of alerts making root cause analysis slow. Nancy’s team came in and implemented automatic correlation that minimized the noise to just the important alerts. This has significantly reduced our MTTM (mean time to mitigate), which improves our customer NPS.”

– **Gabriel Satterlee - Director of Engineering Operations**

This process was the beta, or proof-of-concept (PoC), and was proven to work. Based on these results, we began to push correlation activities out to our other supported products and applications. When we trended the data, post implementation, we noted that there was significant reduction in both MTTT (mean-time-to-triage) and MTTM (mean-time-to-mitigate).

Part of the overall reduction effort involved reviewing the alerts we had in place. By reviewing these alerts, we were able to understand what was actionable, as opposed to a warning or for awareness. We developed a better understanding about how to adjust thresholds where it made

sense, and anticipate what related events could be correlated. We were effectively able to reduce the noise and focus on hidden problem areas via deep dive investigations.

Most of the automation efforts involved escalation reduction, by looking at three factors:

1- Are these escalations actionable or not? If not, we can adjust thresholds to ensure actionable next time or remove them entirely.

2- Reduction of duplicate escalations/tickets: Do we need 5 tickets for the same issue? Typically, that answer is no. Our approach was to merge similar alerts into the same ticket to reduce toil and escalate a single time for a single issue.

3 – Using problem management to focus and deep-dive on the top “repeat offenders.” This helped to prioritize Engineering/Developer teams around specific issues to improve customer experience and reduce load on our teams.

These three factors were key in both deprecating our offshore resources and improving quality of life (work/life balance) for our teams.

~Brian Seeley, Sr. Manager Production Support

As this process unfolded, we realized that we were able to automate a lot of the manual work – meaning that our Tier 1, serving as a catch-and-dispatch team, was not required to support the internal product team, because we were directly assigning to the responsible resolver teams. Those that were not directly assigned to an engineering team went directly to our Tier 2 Triage desk for a more technical review. In some instances, we were even able to move resolution steps all the way to Tier 1 customer Care and, eventually, a customer-facing application, so that they could “self serve” for issue resolution, without having to call us. More on this later.

We did still have the challenge of ticket entry. We lacked a means for our field teams or other business partners to report an issue to us, beyond email. The process involved emailing our Tier 1 team, which would acknowledge the request, create the ticket and assign it to our Tier 2 team for initial review, while pulling of necessary log files. After Tier 2 completed its work, it would be assigned to an incident management team or engineering team, depending on the triage steps. After the Incident Management or Engineering teams completed their piece, the ticket went back to Tier 1, to follow up with the issue reporter. The number of exchanges between assignment groups caused a lag in final root cause or mitigation of the issue.

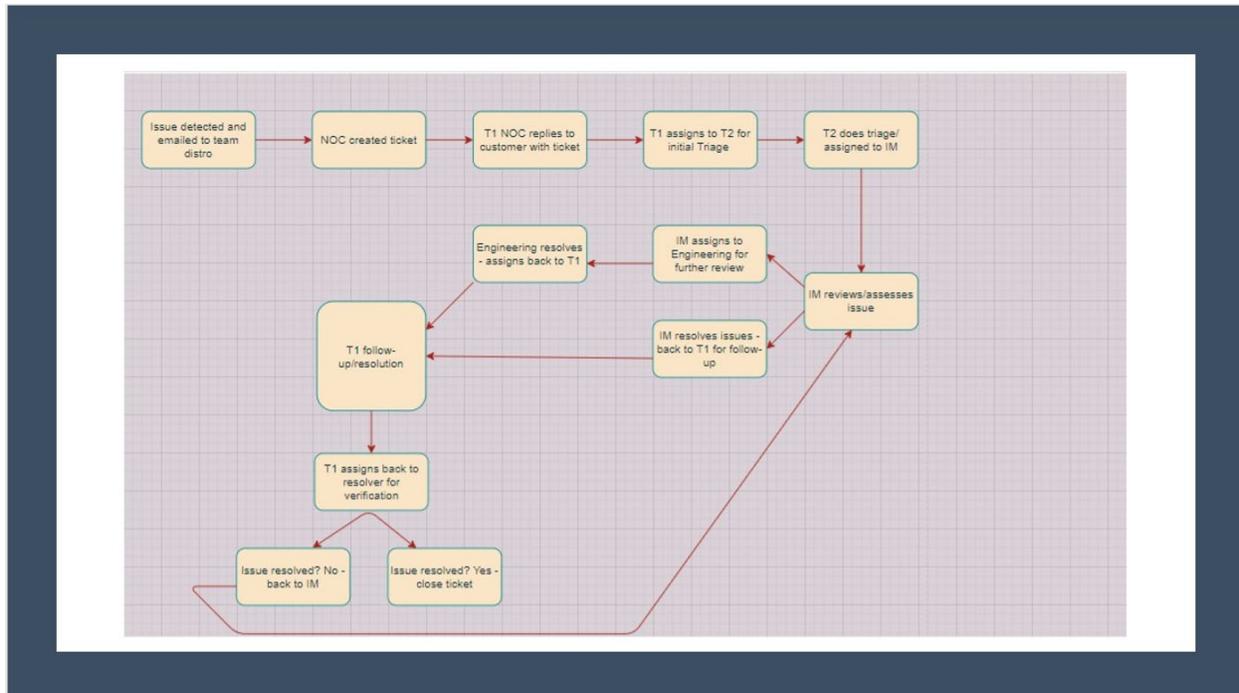


Figure 8- Pre-portal Workflow

Working with our ticketing tools team, we were able to provide requirements for a ticket portal that anyone with the organization could access. This allowed users to directly submit their issues and receive a ticket number at time of entry. We also allowed a user reporting an issue to escalate to the on-call resource, which is especially important after hours or on weekends. **Note, we did include a disclaimer to “use with discretion,” as it will generate a page to the on-call resource’s phone.*

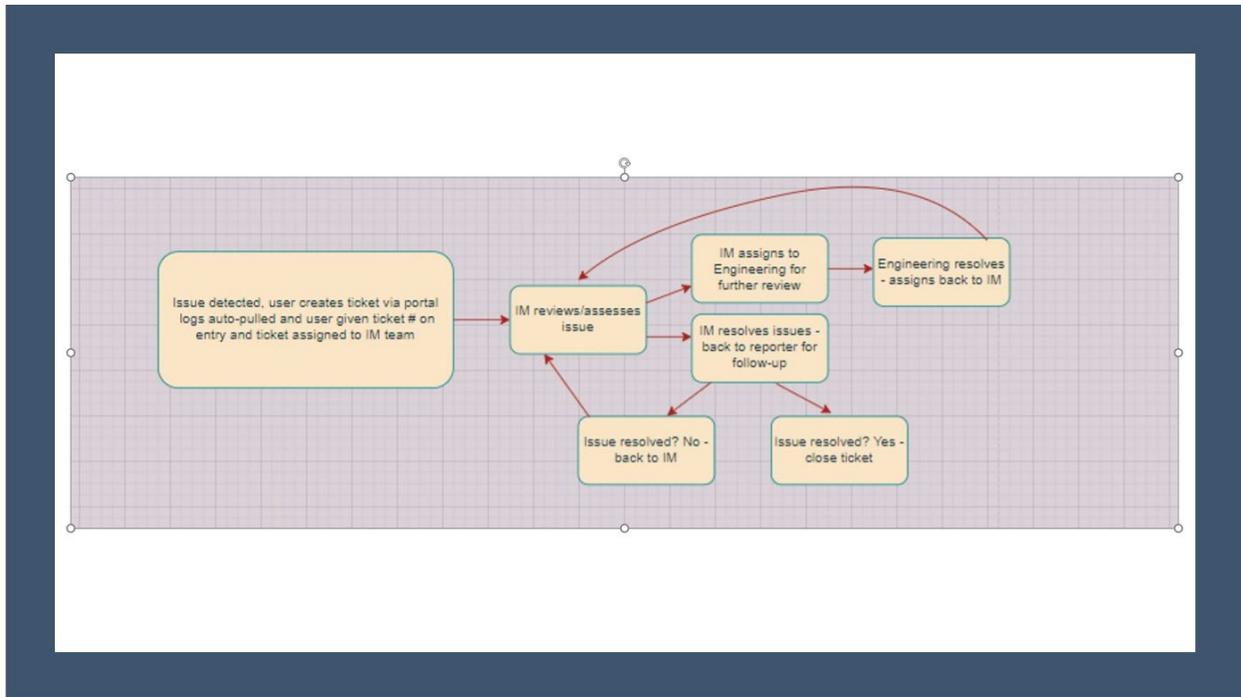
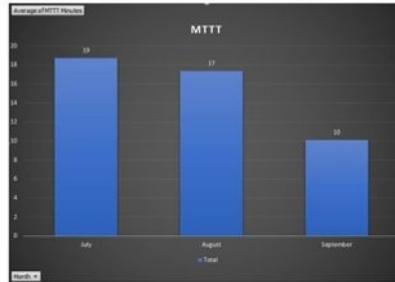


Figure 9- Post-portal Workflow

The portal allowed us to simplify and streamline the ticket process, lessen the number of reassignments, and reduce the MTTM. In addition, the portal and auto-assignment work, we were able to automate pulling device logs when needed (typically a single-customer reported issue). By automating the log pulls, we were able to demonstrate a fair reduction in MTTT and MTTM. We also saved upwards of 2,000 hours (about 2.5 months) of manual labor, and automated most of the Tier 2 work, increasing the opportunity for cost savings.

Mean Time to Triage:

We see an approximate 8-minute decrease for the month of September when comparing to July & August.



Mean Time to Mitigate:

We see an approximate 30-minute decrease for the month of September when comparing to July & August.

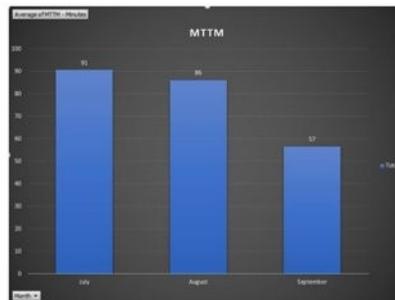


Figure 10 - Mean Time Graphs

Previously, we mentioned tool-sharing and being able to get fixes out to Tier 1 care agents, and eventually directly into our consumers’ hands, via an app. When we subsequently analyzed volume, we started seeing trends. What we noticed was that two API (Application Program Interface) functions accounted for most actions taken to fix a customer’s escalated issue. This required them to call Tier 1 support, which then escalated to the advanced team, which created a ticket for us to email to the Tier 2 team. That is a long way around, when the task is to quickly escalate!

We worked with our development engineering team to safeguard the API and make it accessible to our Production Support team. We validated the functionality of the tool into which these APIs were integrated. Once we signed off on it, we opened it up to the advanced care team. This prevented escalations to us and our engineering partners, but still created an escalation to the advanced care team. It saved mitigation time for the customer -- and we knew we could do better. Comcast uses ITGs (Interactive Troubleshooting Guides) for our Tier 1 agents, so we worked to get the correct LOQ (line of questioning) into the ITG, with a wrapper around the APIs so that Tier 1 agents could execute and resolve the customer’s issue at point-of-contact, by clicking a link. This readily reduced the mean time to resolve as well as escalations – but what if we could help the customer and prevent the call altogether? By leveraging our existing technology and available customer-facing applications, we did just that. We made these functions available directly to the customer in the Xfinity App.

9. Conclusion

A careful review and analysis of your network, systems and applications helps identify areas of sameness, best practices, and resultant opportunities for process alignment and automation. When these opportunities are missed, it can slow the process of automation, potentially causing re-work. Categorizing your findings helps define your proof of concept and identify the low hanging fruit (the quick wins). We found it helpful to write out a problem statement, including an anticipated solution or expected result. From there, determinations can be made about what tools to leverage, to get from problem to solution, and to start defining the “How” of automation opportunities.

In reviewing the automation journey, we categorized opportunities into three “buckets.” The first was the “low-hanging fruit” / quick wins, which in our case was the ticket auto assignment. The second was event correlation, which involved a comparatively medium level of effort. We had the foundation already available but needed to ensure that we had proper mappings and service flows in place before implementation. The third was the more difficult, longer-range automation plan of getting resolutions pushed out customer-facing tools, like our customer care app, Xfinity Assistant. This was a multi-phased approach that took a lot of coordination with various teams.

What drove our decisions was the result of extensive, targeted data analysis, which identified the redundant task-type work. If automation is on your to-do list, start by looking for repetition – what tasks are team members are being asked to perform over and over? This also helps to prioritize your automated efforts. By analyzing our email platform, for instance, we were able to determine that there were two main functions, or API calls, that were leveraged to resolve in the range of 90% of customer reported issues. This was a separate opportunity, of the windfall variety, that was identified while we were reviewing options for event correlation.

One consistent theme that emerged throughout all our automation efforts was that it must be measurable. We have the data to baseline where we started, and can continue to measure results to prove (and occasionally disprove) the value of the work being done. Are we seeing the benefits of this effort, as we anticipated? If yes, proceed. If no, re-evaluate: Either we are missing a key component, or this really does not have the ROI (return on investment) we originally envisioned. In which case, it may make sense to deprioritize this automation task and move to something that will have greater impact.

Keep in mind: Data can come from a lot of sources. It can also be a combination of different source types that can give you what you need to measure performance and results accurately. Sometimes, you’ll find that you have this great idea, but no way to prove it because the data does not exist. That does not mean you throw away your idea! Reframe the question. Work to identify, define, and implement the right logging to build your data. It is likely a gap that needs to be addressed.

Last not least: Stay agile. Flexibility and adaptability should be the forefront of any effort. Remember to identify the big wins as well as the low-hanging fruit. Ensure that your data tells the story and paints your vision, which will help you to garner the support of your teams for buy-in. Because ultimately, talking about automation is great, and we can sing its praises -- but without stakeholder and customer buy-in, it is not going anywhere. It is extremely difficult to enact change. Most employees, team members and even family members are resistant to changing the way they do their daily tasks. By understanding other teams' processes and procedures, you can highlight the benefits they will see by implementing the

suggested methods. You will also find that they have critical processes that you may want to leverage into your playbook – which brings us back to pollinating best practices throughout the organization.

Abbreviations

API	Application Programming Interface
HFC	Hybrid Fiber Coax
ITG	Interactive Troubleshooting Guide
LOQ	Line of Questioning
MSO	Multi-Soaking Outage
MTTD	Mean-Time-To-Detect
MTTM	Mean-Time-To-Mitigate
MTTT	Mean-Time-To-Triage
NTF	No Trouble Found
OPV	Outage Pre-Verify
VM	Virtual Machine
XOC	eXcellence in Operation Centers