VIRTUAL EXPERIENCE
OCTOBER 11-14

2021 Fall Technical Forum
SCTE • NCTA • CABLELABS

The Zen of Ticketing

Operational Transformation

A Technical Paper prepared for SCTE by

Melissa Wood
Sr. Director Reliability Engineering, Program Management
Comcast
183 Inverness Drive West, Englewood, Colorado
719-491-0457
Melissa_wood@cable.comcast.com

Table of Contents

Title **Page Number**

1. Introduction ... 3
2. Ticketing Infrastructures: The Backstory ... 3
3. The Challenges of Managing Outages and Network Upgrades ... 4
4. Transparent Ticketing Processes Create Reliability ... 5
   - 4.1. Analysis Overview ... 5
   - 4.2. Future State Network Enterprise Architecture ... 9
   - 4.3. Implementation Plan ... 10
5. Agile Solutions With Continuous Learning are the Way ... 12
   - 5.1. Lessons Learned and Pro Tips ... 13
   - 5.2. Success Stories ... 14
6. The Future of Ticketing ... 15
7. Conclusion ... 15

Abbreviations ... 16

Bibliography & References ... 16

List of Figures

Title **Page Number**

Figure 1 – ITIL Core Components ... 6
Figure 2 – SRE Key Principles ... 7
Figure 3 – Site Reliability Engineer Responsibilities ... 8
Figure 4 – ITIL Lifecycle Architecture Support ... 10
Figure 5 – Implementation Team Structure ... 11
Figure 6 - Enterprise Architecture Implementation Plan ... 12
Figure 7 - The Three Dimensions of a Learning Culture ... 13

List of Tables

Title **Page Number**

Table 1 – The Zen of Ticketing Platform Analysis ... 9

1. Introduction

Enterprise ticketing systems are among the vital subsystems that "keep the lights on" for broadband service providers – or, more accurately, keep service-related lights off. Service-related tickets, not unlike the complex and interwoven landscape of "tools," in general, are foundational for identifying, tracking, and resolving internal and external operations, from scheduled maintenance to incident handling. For those reasons, ticketing systems are often as layered and difficult to detach as multiple layers of wallpaper on a wall -- they got that way through decades of growth, patches, ownership changes, and management shifts (e.g. from centralized to decentralized and back again).

As the title of this paper indicates, this paper is about how to get to "zen" when it comes to incident and trouble ticketing. We define zen as a general state of enlightenment attained by summing the parts, which in this case, are the many different ticketing systems acquired through decades of individual acquisitions that helped defined our industry's consolidation.

In this paper, we will share how and why Comcast pivoted toward a unified ticketing system -- including what constitutes realistic expectations, lessons learned, useful metrics, and best practices. We will characterize what it takes to unify multiple ticketing systems, so as to attain the benefits that come from operational scale.

2. Ticketing Infrastructures: The Backstory

Comcast provides multiple services to customers to help keep them connected, everywhere, all the time, across a large geographical footprint. From traditional cable television, to Internet, phone, security, and wireless, Comcast and the cable industry deliver products every day that consumers expect to 1) be there when needed, which is all the time, and 2) evolve gracefully.

Overall consumer desires for service availability and speed puts extra emphasis on our engineering and operations teams to execute reliability measures in ways that get a "thumbs up" from consumers every time. Moreover, managing changes to the network production environment, so as not to cause self-inflicted outages, is part of what prompted Comcast to move towards a unified and enterprise-wide ticketing infrastructure.

Why does this matter now? The complexity of the network topology, coupled with heightened consumer expectations, requires a scalable solution. We can no longer "throw more people at it," when it means trying to manually understand all the many cross-departmental tickets and alarms, all firing inside an increasingly complex network topology. The answers require people and technology.

Like many of you reading this, Comcast "got here" through multiple acquisitions, as a result of a heavy few decades of industrial consolidation. This coupled with understandable and necessary localized network requirements at the regional and divisional levels, produced differing hardware, software, and – tantamount to this discussion – ticketing systems, especially as it relates to managing outages and network upgrades. As a direct result, processes and tools vary across our different engineering and operations teams.

Collaborating during an upgrade or outage event demands teams chat on different platforms, and to document ticket information in multiple places to keep stakeholders informed. These variations cause unintended delays to resolve the customer's impairment, resulting in their inability to use services they pay for. The ability to scale and reliably manage the network is more important than it has ever been.

Basically, to provide new and highly reliable services to customers, at scale, requires pioneering towards one network, end-to-end -- including ticketing -- across people, processes, and technologies.

Moreover, our employees provide great feedback on a regular basis about on-the-job processes that are difficult or could be approved upon. Now is the time to make those processes straightforward with the same execution regardless of the geographic location. Building a unified enterprise ticketing infrastructure sets the tone and removes assumptions, which gets us closer to full operational transparency.

3. The Challenges of Managing Outages and Network Upgrades

Tools and processes across Comcast vary when it comes to managing outages and performing network upgrades. Any service provider that grew its geographic footprint by acquisition, then resolved to consolidate and cluster its holdings, invariably faces an avalanche of tools, like ticketing systems, that are used regularly, thus are familiar, yet they all do essentially the same thing. Our internal analysis, for instance, revealed numerous siloed applications that existed to support teams doing incident and change management. Such replication of records adds risk to delivering reliable services -- risks that are potentially service-impacting, thus customer-impacting. Like all service-minded broadband providers, we continually upgrade our network to provide more products and services, with the very specific intent to be "always on." This good work can be overlooked, especially if an outage accompanies an otherwise proactive change.

Data or configuration items (CI) tend to be stored in many different tools and can lack a standard way of providing relationships between them. A CI represents service components, infrastructure elements that need to be managed for successful delivery of services. An example is network components, such as, routers, hub, gateway, etc. Technologies vary greatly from group to group, when it comes to how hardware and software is monitored. Teams tend to work outages from the same CI, producing up to 20 tickets for a single impact, which also produces multiple instances of repeat communications.

Impacts in the network drive impairments to applications that both customers and employees use. When individual teams are working in a silo to conduct incident and change management, the number of incident tickets are inflated, and duplicate technical and executive communications are distributed. Change tickets themselves can result in an outage, due to not having robust conflict assessment across the company. While we are in a far better place than we were 5 to 10 years ago, when it comes to "swivel chair" operations, engineers still spend time manually looking up needed information to solve an outage, having to swivel between tools to manually enter data needed either in an outage or change ticket.

Another contributing factor is alarm correlation across topologies and different activities. When alarm correlation is weak, it is often because root cause analysis is often only done on "severity one and two" tickets, which only equates to about a quarter of the workload. Lower-level severity tickets are created only to form documentation to identify patterns and trends. "Parent" and "child" tickets aren't always properly linked, which puts a weak spot again on correlation. The amount of manual work performed by engineering and operations teams can delay outage mitigation, degrade the customer experience, and limit the ability to scale needed changes within the network.

The guiding principle towards any major overhaul to how development and operations teams work together is to create a vision with shared visibility, communication and collaboration that embraces working together in a blameless culture to drive continuous improvements, always.

4. Transparent Ticketing Processes Create Reliability

Developing a transparent ticketing process for higher reliability is easier said than done! Assessing the nuts and bolts of how your organization is operating today to the desired future state will create a shared vision for alignment utilizing best practices others have identified.

A reliability vision spans people, process, and technology. From a people perspective, there is a need to have clear, accurate and timely communications to the appropriate technical, executive and customer audiences. Processes should establish a robust, data-driven culture to drive resource allocation appropriately. The technology should provide a framework to build automation that decreases outage impacts and reduces unnecessary and especially duplicative labor.

There are several steps outlined in this section that helped to organize the Comcast journey. These steps provided a framework that guided teams through design, to the transition into the zen of ticketing, and eventually to operationalizing ongoing practices.

4.1. Analysis Overview

The first step in this zen of ticketing journey was to identify a solid framework (or two!) to manage the enterprise architecture. The three most common that engineering and information technology (IT) leaders lean towards to accelerate automation of things are information technology infrastructure library (ITIL), site reliability engineering (SRE) and development and operations (DevOps).

Our approach involved blended principles from SRE and ITIL to drive innovation through an agile continuous learning model. ITIL has provided common best practices helping across many operational teams, while SRE brings to light self-service tools and automation scripts. The models together increase reliability and performance of applications and services as manual work decreases. With a constant eye toward automation of manual processes brings the ability to auto-mitigate, remediate, and removes manual labor.

First, we will address the ITIL model. The ITIL is a five-step process to measure, report, plan and implement quality improvement cycles for ticketing operations. Its core components, including Service Strategy, Service Design, Service Transition and Service Operation, provide the framework for continual service improvement. The ITIL service management lifecycle provides guidance for implementation of a configuration management database (CMDB), event, incident and change management with policies, guidelines, and streamlined processes. The is extremely important, as it provides the capability to build physical and software configurations items into relationships that serve as the base for incident and change tickets. It also sets the foundation for automation, and aids in understanding customer impact. Figure 1 shows the ITIL core components (Ghadi, 2011).

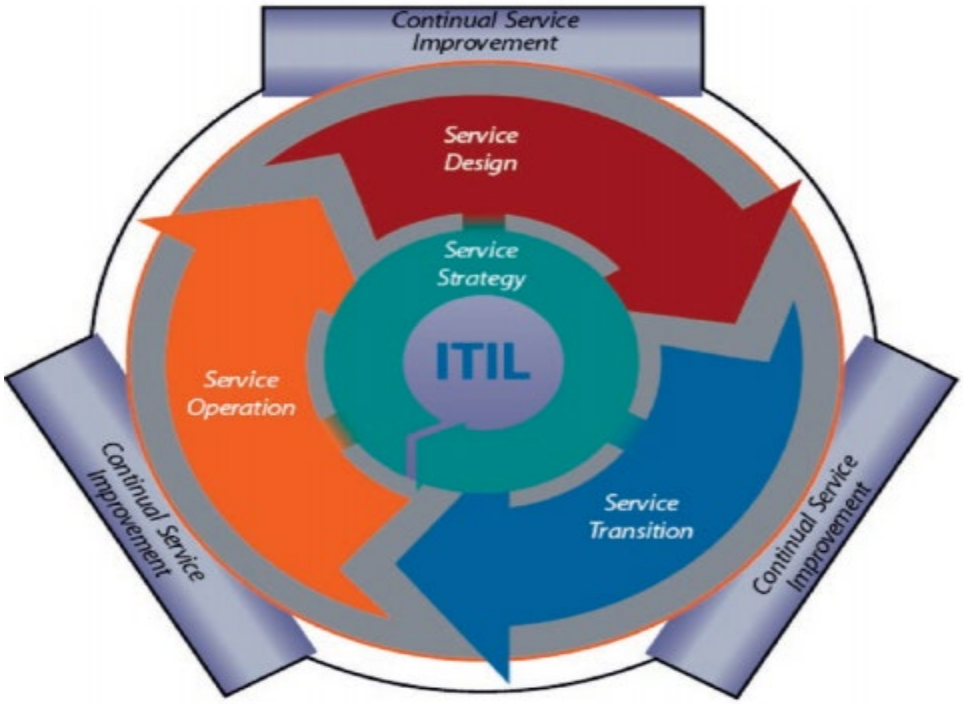

Figure 1 – ITIL Core Components

The ITIL components guided us through the common network enterprise architecture layers of business, data, solution, and technology. The focus on the four common network enterprise architecture layers was intended to reduce duplication, and to set the stage for automating certain repeatable tasks. The layers, informed by the focus needed to achieve a single ticketing platform, included:

- Business architecture or processes to define the strategy for moving people, process and tools to a single ticketing platform. Such as "governance" to inform key business processes.
- Data will include a standardized CMDB to identify physical assets, with relationships that tied to incident, change and problem management tickets.
- A solutions architecture to provide a blueprint of the end-to-end ecosystem, especially with regard to how applications will connect and get deployed.
- A technology architecture to assess what software is appropriate to support the evolution of a single ticketing platform. This view defines applications, databases and how they bridge together.

Each layer provided a different view to achieve the desired ITIL enterprise architecture. The viewpoints consisted of the following components:

- Business
    - A detailed vision to achieve best practices to transform the enterprise to a single ticketing platform
    - An ITIL Service Strategy to guide the design, development, and implementation of service management lifecycle with policies, guidelines, and processes

  - A review of current practices to understand "must have" end user requirements to align to future state
- Data
  - CMDB model with relationships diagramed
  - Schema of Incident, Change and Problem workflows
- Solution
  - Ecosystem design to map out process and software connections
  - Defined service level agreements supporting Incident, Change and Problem processes
- Technology
  - Identification of tools and service delivery platform

The second model we utilized is at the heart and soul of Google as they discovered a need to manage risk and growth called "Site Reliability Engineering" or SRE. SRE goes down the path of having software engineers that have the drive and ability to automate across complex architectures. Scale and reliability are created to manage higher volume of changes to your network with this methodology.

Figure 2 – SRE Key Principles

As shown in Figure 2, SRE drives resiliency into infrastructure and workflows that provide automatic responses with code to what humans have previously done (Site Reliability Engineering, 2020). This practice can transfer operational work into development tasks. As with the ITIL framework, SRE has five base elements to guide towards a software engineering model:

- Business communication for clear alignment on the definition of reliability with service levels that have associated impact.
- Architecture creation of scalable systems that have resilience to reduce outages.
- Build and run using automated toolchains for provisioning and deploying code is the foundation of automating manual work.
- Operating and monitoring aspect to measure everything that matters via service level indicators (SLI) and service level objectives (SLO) to assist with business impact.

- The SRE culture is set up to define a mix around 40% software and 60% administration system capabilities, with each SRE having accountable SLOs.

Implementing an SRE environment provided a higher ability to complete errorless transactions with proactive monitoring that frees the team to work on structural improvements at a more rapid pace. SRE provides not only definition for how to improve reliability, but also provides engineers with the opportunity to build and design instead of just putting out fires. Google has taken lead with the SRE role and methodology to foster sustainability and operational resiliency of all digital assets. Figure 3 provides a sample of what a SRE engineer focuses on (Feoktistov, 2021).

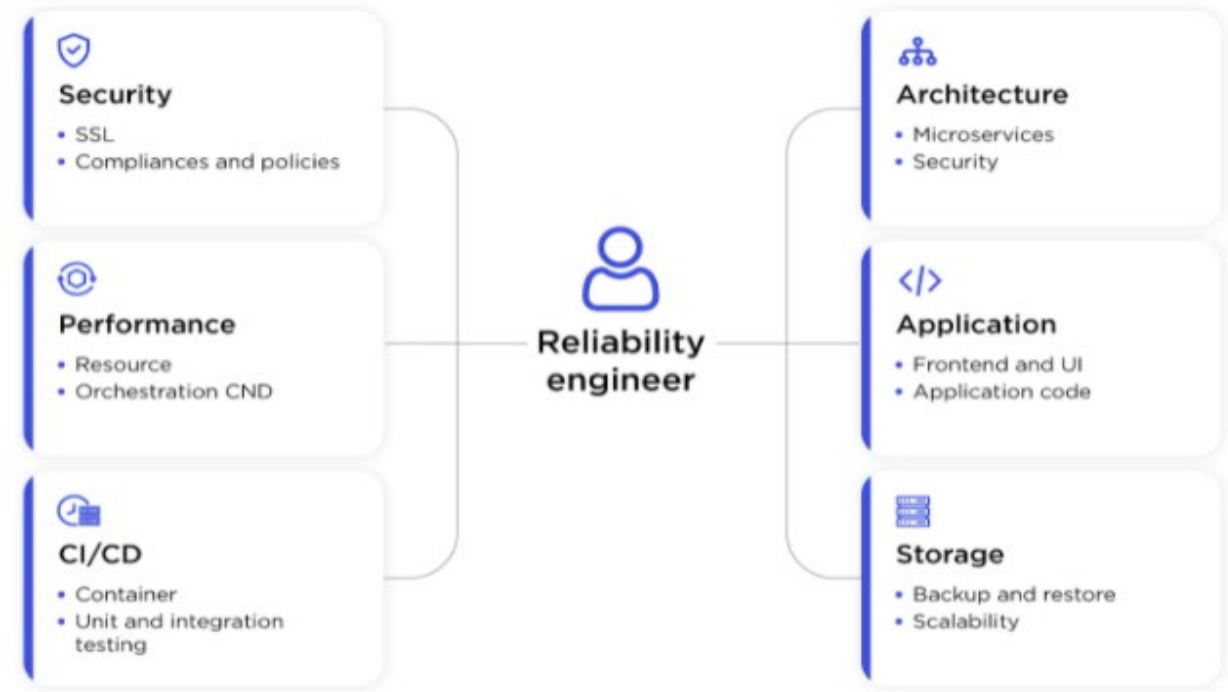

Figure 3 – Site Reliability Engineer Responsibilities

Comcast benefited from a blend of ITIL and SRE to address processes across the organization, as well as the complexity of the CI that existed, with a focus on eliminating manual work. The next step involved mapping out the benefits of this approach, in order to address the question of why any company would break down their current ticketing infrastructure and drive their teams into a single platform. Streamlining incident and change tickets to lower the amount of manual work, improving communications and building an ecosystem to remove toil to help set the stage for scalability are but a few answers to that question. To get to where you are going, you must understand where you have been. Table 1 defines the who, what, why, when, how and where of mapping the current state of ticketing to the future zen state of ticketing. This exercise is intended to understand the different audiences and how they will be impacted.

Table 1 – The Zen of Ticketing Platform Analysis

Classification Names / Audience Perspective	**Who** People, resources	**What** Entities involved in each perspective	**Why** Goals, objectives & business plan	**When** Time and performance criteria	**How** Functions in each perspective	**Where** Locations and interconnections within the enterprise	Classification Names / Model Names
Executive Perspective	List of all teams doing Incident and Change tickets.	List of what each team is responsible for, inventory list and best practices for Incident and Change tickets.	List of benefits for moving to a single ticketing platform.	SLAs for managing the Incident lifecycle from open to mitigation to close; Change ticket lifecycle.	List current & future state to identify transformations to include governance process for future state.	Map of geographical locations to show how they are interconnected.	**Scope Contexts**
Business Management Perspective	Define roles and responsibilities across teams.	Process flow diagrams showing inventory across processes, to include CMDB	Remove duplicate tickets; Create change risk model and conflict assessment.	Determine time needed to get to future state.	Diagram inputs and outputs across process transformations.	Define distribution.	**Business Contexts**
Architect Perspective	Understand system roles to define CMDB for inventory to feed Incident and Change tickets.	Schematic showing inventory across processes, to include CMDB.	Diagram showing system end state CMDB lifecycle.	System timing representation.	System process representation.	System locations.	**System Logic**
Director Perspective	Technology role diagram.	Inventory configuration.	Software selection.	Timing to implement technology.	Software process representation.	Technology locations.	**Technology Physics**
Engineer Perspective	Tool and security roles.	Data definition	Rules for software implementation.	Timing to implement technology.	Process configuration.	Distribution configuration.	**Tool components**
Enterprise perspective	Organization	Data	Strategy	Schedule	Function	Network	

4.2. Future State Network Enterprise Architecture

As stated, ITIL provides a recipe for reliability practices. Figure 4 shows the main areas of focus to improve incident and change management processes. Comcast chose a third-party platform with plugins to align with overall business strategies, including financial planning, applications testing, agile development, and project management, to name a few. The Service Strategy, Service Design, Service Transition and Service Operation guided the team toward agreed-upon implementation plans. An internal SRE team was created to configure the platform towards the aligned future state with a focus in the following areas (IT Service Management, 2019):

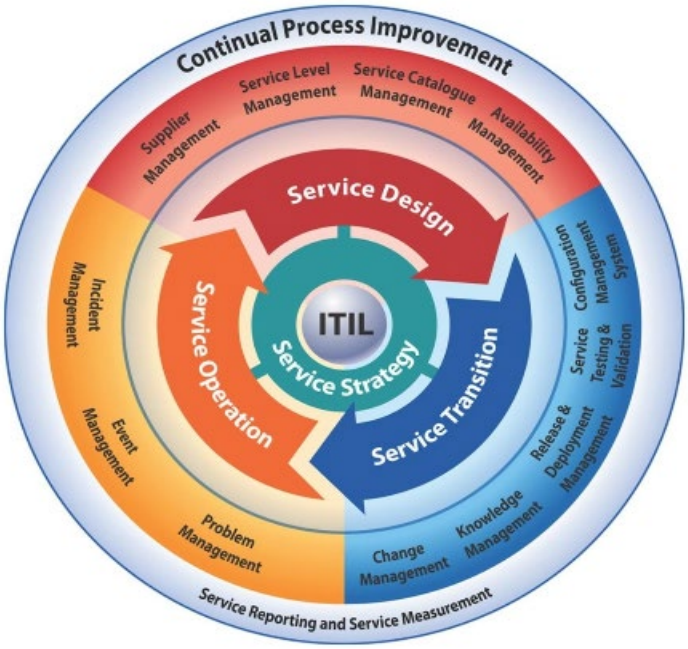

Figure 4 – ITIL Lifecycle Architecture Support

- **Configuration Management Database (CMDB)** – provides the capability to build physical and software configuration information into and between relationships that will be the base of incident and change tickets. This sets the foundation for automation and aids in understanding impact.
- **Incident Management** – provides a single platform for cross-departmental teams to work in unison on outages, with a focus on improving time to engage, time to mitigate and time to restore an outage with strong reporting.
- **On-Call Scheduling** – provides automated escalation from an incident ticket for fix agents to get engaged on an outage.
- **Problem Management –** provides the ability to perform trend analysis, within the same single platform, while tracking internal and third-party vendor action items
- **Change Management** – provides a single platform with a systematic, trackable and controlled approach to the lifecycle of all changes, enterprise-wide, to understand conflicts and customer impact.

4.3. Implementation Plan

Migrating to a zen of ticketing state requires a lifecycle strategy. Collaboration between the team managing the business change, and stakeholders of all impacted teams is key. The zen of ticketing strategy, goals and teams were documented and approved by senior leadership to ensure alignment with the future state. Figure 5 provides an example team structure to build a focused team.

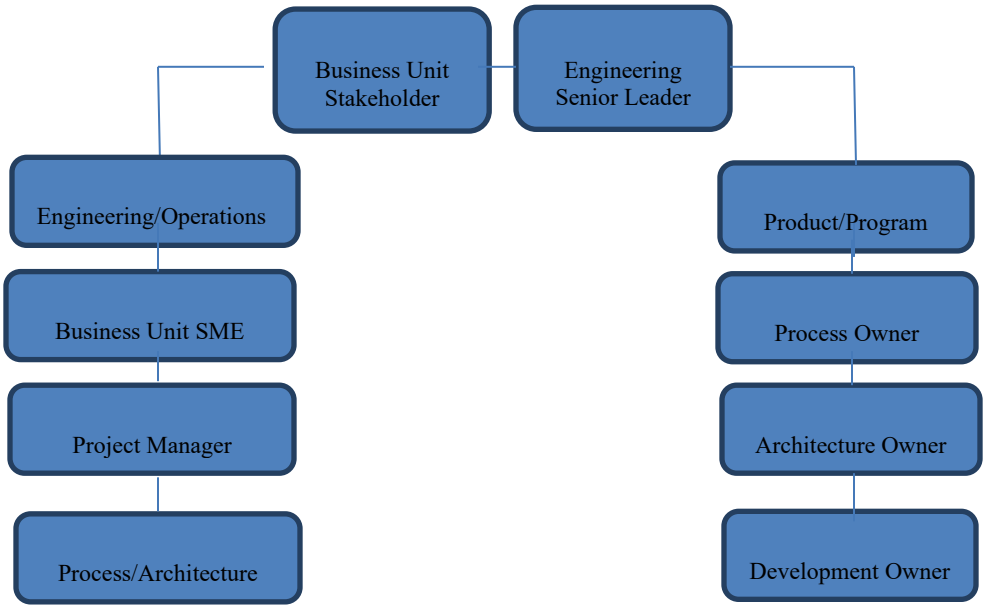

Figure 5 – Implementation Team Structure

Implementing the role of a "product owner" provided a liaison structure across multiple teams to learn and decide how to move through the lifecycle implementation plan. As the discovery process took place, the product owner facilitated the conversation to understand each team's current as-is state, mapping to the future state. With the support of senior leaders, the business goals identified a future visionary state that included the benefits of implementing a new architecture. One might call this a product roadmap and/or charter that is then socialized with teams during a kickoff session. The product manager worked with stakeholders to identify what the "day one" experience would look like, and the architecture team gathered the information it needed to draft different viewpoints to illustrate the alignment.

As the experts on the new technology, the process team mapped current workstream processes to the future state, being careful to identify areas where people would need to change what they were doing. Once the new architecture is approved, the product team will create and submit "stories" for the development team. A story in this context describes what the business is aiming to achieve with incremental code development. The number of stories identified will determine length time needed for development.

The development methodology is necessarily agile and iterative. As functionality is ready to be reviewed, visual demos and user acceptance testing will commence. As each functionality piece is finalized in the lower-level environment, process documentation is created, which will also support training efforts. Training needs may vary by team; therefore, this segment is handled uniquely to ensure the "people readiness" aspects cover awareness, desire, knowledge, and reinforcement. The overarching goal is that teams understand the "why" behind the change. Figure 6 displays the milestone implementation steps for the focused areas each team will need to accomplish.

	Workstreams	Discovery (3 months)			Design (3 months)				Development (3 months)			Onboarding (3 months)		Gating Dependency	Risks	Mitigation
		Kickoff	CMDB Current	Process/ Tools Current	Access Setup	CMDB Future	Process/ Metrics Align	Prod Day 1 New	Req Stories	Arch	Dev UAT	War-game	Training			
Business/People	Vision, Goals & Benefits													-	-	-
	ITIL Service Strategy													-	-	-
Data/Process	Configuration Management Database													-	-	-
	Incident Management													-	-	-
	On-Call Scheduling													-	-	-
	Problem Management													-	-	-
	Change Management													-	-	-
Technology	Core Platform													-	-	-
	Integrations													-	-	-
	Middleware													-	-	-

LEGEND: On track; At risk; Blocked; Completed; Not Started

Figure 6 - Enterprise Architecture Implementation Plan

The mission-critical business drivers encompass improvements that are required in the data and process sections of the workstreams. This, along with developing the overall technology ecosystem, will drive duplicate tickets to one in the single zen platform. Repeatable steps will be identified through discovery, with a focus on automating where it makes sense. The implementation lifecycle can vary depending on the complexity of the network topology and availability of source data. Designing a governance process will help manage the aligned continuous improvements, timeline, and resources required. Finally, regular communications are critical to keep the program team engaged along with leadership across the enterprise. Long term programs can easily succumb to distractions if there is not discipline around program and communication management dwindles.

5. Agile Solutions With Continuous Learning are the Way

Organizations that adopt learning as a practice increase the chances of continuous improvement, while promoting a culture of knowledge and innovation. Leading with a mantra of transparency in all operational and engineering practices also helps to cultivate a blame-free culture. Agility sets boundaries for how to work, in preparation for being able to pivot to new processes and technology. Those pivots may arise because of business needs or competitive threats; in all cases, the customer always plays the primary role. As a result, team environments that achieve shared milestone objectives are more productive, and more emotionally satisfied with their work. A scaled agile framework provides three critical dimensions to help build a learning culture, as shown in Figure 7 (Continuous Learning Culture, 2021). All three dimensions speak to employees at every level to foster growth and transformation.

Figure 7 - The Three Dimensions of a Learning Culture

5.1. Lessons Learned and Pro Tips

As we continue our journey towards the zen of ticketing, we identify a few pitfalls to help our readers avoid those things that can disaffect their single zen ticketing aspirations.

The first lesson learned involves the objective, the goal, and the timeline for a single zen ticketing initiative. Most cable operators already have a ticketing system (or 20!) in place that may have been created over the past 15-20 years. Be sensitive to the time and people that it took to build what worked well until now. Set smaller and realistic timelines that create incremental improvements over time. A governance process will help keep those involved grounded as to what the next quarter's commitments will bring.

A second lesson learned involved the CMDB. As described earlier, the CMDB provides the capability to build physical and software configurations items into relationships that serve as the base for incident and change tickets. It also sets the foundation for automation, and aids in understanding impact.

"Do not underestimate the value in moving towards a single solution that can provide automation and correlation of events, so that engineers can focus on the problem instead of being ticket jockeys," noted Rich Massi, VP, Reliability Engineering Residential Products and Services, who shared his viewpoints during an interview. Build a process to stay current, because the CMDB is never done. Also important is how and who best maintains the CMDB's governance and data quality processes.

Automation is a key aspect of moving into a single zen ticketing platform. Covet all insights about how to perform a task, should automation break or go down -- it happens. Ensure everyone understands manual processes from the beginning, before automating tasks. (You can thank us later for this one!)

Call it what it is. The bigger the organization, the higher the likelihood that big software initiatives like ticketing will come with a fancy name, a logo, stickers and more. Just call it "ticketing"! This will reduce

financial and capital expenditure discussions to come, allowing you to focus on the incremental improvements.

Product and process are alignment upfront. The process should drive the implementation as opposed to the tool trying to drive the process. It is important that there is alignment on the process before the tool is adapted. Otherwise, just another problem is created when multiple groups are trying to modify the tool in competing ways.

Lastly, and speaking of the finances: think about the big picture. Larger companies going through large transformations like moving towards a single ticketing platform should anticipate funding requirements. Considerations like vendor support versus in-house work, i.e., "buy versus build," will depend on what skill competencies exist to start such a major program. A multi-year program, unified ticketing is best approached as an endeavor that is forecasted and budgeted to cover operational and capital expenditures.

5.2. Success Stories

We're starting to gather the success stories that come with moving our teams onto a single ticketing platform. We're not all the way there yet – about halfway, at the time this paper was written. Yet so many wins have been and have yet to be celebrated! Results so far clearly indicate that moving incident and change tickets into the single zen ticketing platform is reducing "swivel" and the time it takes to manage those processes.

The voice service engineering and operations team were the first to embark, diving headfirst into the single ticketing platform. Frances Augustine, Executive Director Reliability Engineering, led the vision to reduce the number of alerts engineers were managing. She was relentless that "eyes on glass" was going to become a thing of past, and that managing with a proactive mindset would be a much healthier work environment. Through the automation associated with a single ticketing system supported by a mature CMDB, the "Voice team" was able to do just that and move from a reactive to proactive operation, focusing on a sustainable level of logged alerts to tackle before an outage occurs. This includes internal alerts, to ensure voice services are always on, to providing automated reporting to the Federal Communications Commission (FCC) for its required outage notifications. This automation enabled engineers to pivot from being reactive to being proactive, finding issues from logs before they became outages. It also enabled a reduction of Tier 1 vendor support.

Another early adopter was in our residential services and products team. Early on, this team aligned with the single ticketing platform, spending hours to build out its CMDB so that all the appropriate software, hardware and respective locations were identified, connected, and correlated when doing incident and change tickets. The extensive time spent to build (while knowing that CMDB work is never really done; see section 5.1) also enabled automation of incident and change ticketing. As a result, the team was able to eliminate an entire vendor team that had provided manual triage support. This team not only tackled the single tools aspect but is also now ensuring that all teams performing incident management are doing it the same way, so as to harvest additional business improvements.

As part of the journey to adapt to a single ticketing platform, having key performance indicators (KPI) per team, across a collective set of teams, will help measure the success in a few sample areas:

1. Mean time to mitigate and resolve incidents
2. Change success percentage (number of changes executed/not executed flawlessly)
3. Number of changes automated
4. Number of incidents created automatically from event management logging and correlation
5. Percentage of time team spends in reactive versus proactive tasks over X Time

6. The Future of Ticketing

The future of the ticketing, zen or otherwise, is really all about scalability. It's about acknowledging industry consolidation, building smart network platforms, and utilizing intelligence in the right way to ticket the things that matter. Tickets are not necessary for every logged event -- but every logged event can provide trend and analysis to move towards managing proactively. Simply put, our networks are smart and intelligent -- but if we're not smart about our networks, we can't let them show their intelligence. Single ticketing is the path to network intelligence.

The nirvana towards moving into a single ticketing platform, besides the obvious benefits of removing unnecessary duplication and effort, is the uplift to overall reliability. "Spend a little bit of time on Incident Management, police Change tickets that fall out from automation, and spend 80% of your time doing problem management to enable even more automation," advises James Manchester, our SVP of Core Platform Technologies.

What's next? Everything that enables reliability to be "always on" and flourishing. As incident and change ticketing puts firm solid practices in place for ticketing, next steps will likely include moving teams into a single problem management space. Doing so will enable further analysis to ensure an outage, if it happens, only happens once.

Cybersecurity is another realm to explore with respect to unified ticketing. Knowing what is happening in the network is always important, and more so in digital times. Professional hackers continue to advance where and how they can breach systems. It's entirely plausible for managed ticketing practices to also advance, to keep would-be interlopers out.

7. Conclusion

Getting to a "zen state" as it relates to change and incident ticketing, or any type of ticketing, is really about "keeping the lights on" and driving reliability as a core service, at a fundamental and strategic level. This is especially true for large companies, like ours, that grew as a result of multiple decades of system acquisitions and geographic clustering, then remained "siloed" through onboarding and corporate integration. As it continues to become mainstream within Comcast, the unified zen of ticketing approach will bring valuable intelligence to understanding topology and automating remedial tasks.

We'll close with this list of tips, when considering a unified ticketing environment where you work:

- Think about a unified enterprise ticketing infrastructure as a way to create alignment on tools, processes, and people.
- Have a vision to steer teams towards.
- Know that transparent ticketing processes create reliability.
- Adopt a framework such as ITIL and/or SRE to design, transition, and operate into a single ticketing platform.
- Analyze the current and future state of tools, processes, and people, to understand what's going on.
- The right team and leadership support structure is crucial for stakeholder support.
- Under-commit and over-deliver wherever possible!
- Go slow to go fast in an agile continuous learning environment.
- Perform retrospectives with a culture of learning; be an advocate for learning.
- Celebrate the successes and continue to look to the future.

Abbreviations

CI	configuration item
CMDB	configuration management database
DevOps	development and operations
FCC	Federal Communications Commission
IT	information technology
ITIL	information technology information library
ITSM	IT service management
KPI	key performance indicator
SLI	service level indicators
SLO	service level objectives
SRE	site reliability engineering

Bibliography & References

Augustine, F. (2021, June 13). Personal interview.

Continuous Learning Culture. Scaled Agile, Inc. (2021, June 24). Retreived from https://www.scaledagileframework.com/continuous-learning-culture/

Connors, B. (2021, June). Company meeting.

Comcast. (2019, April 24). Company. Retrieved from https://corporate.comcast.com/company

Feoktistov, I. (2021). Why and How to Hire a Site Reliability Engineer (SRE). Retrieved from https://relevant.software/blog/hire-site-reliability-engineer/

Ghadi, V. (2011). Adopting ITIL Framework. Retrieved from https://www.happiestminds.com/whitepapers/Adopting-ITIL-Framework.pdf?sid=5086

IT Service Management. (2019). Retrieved from https://docs.servicenow.com/bundle/madrid-it-service-management/page/product/it-service-management/reference/r_ITServiceManagement.html

John P. Zachman. (2019). The Zachman Framework Evolution by John P Zachman. Retrieved from https://www.zachman.com/ea-articles-reference/54-the-zachman-framework-evolution

Manchester, J (2021, June 13). Personal interview.

Massi, R. (2021, June 13). Personal interview.

Minoli, D. (2019). Enterprise Architecture A to Z: Frameworks, Business Process Modeling, SOA, and Infrastructure Technology. Auerbach Publications.

Site Reliability Engineering (SRE) What is it, what are the advantages, and how do I implement it? (2020, March 31). Retrieved from *https://www.linkedin.com/pulse/site-reliability-engineering-sre-what-advantages-how-do-dan-martines/*

Young, C. (2021, June 4). Personal interview.

