



Scaling On-Call in a DevOps Organization

Contents

3	On-Call Fundementals
4 4 5	On-Call Concepts What is an on-call schedule? What is an escalation?
6	The DevOps Movement and Its Crucial Use of On-Call Scheduling
7	On-Call Schedule Patterns for Growing
	Companies
7	One or Multiple People On-Call
8	Multiple Escalations
8	Escalating to a different team
9	Separate nightwatchers
9	Schedules with Overlaps
10	Follow the Sun Schedules for International
	Organizations
11	Advice to Management Teams While Enrolling Changes
	to On-Call Systems
11	Transparency
12	Training
13	Development Duties During On-Call
14	Well-Defined and Fined-Tuned Processes and
	Systems
15	Reviewing On-Call Reports
16	Friendly On-Call Culture

On-Call Fundamentals

Just like emergency rooms require on-call schedules for doctors to handle health emergencies, DevOps teams need on-call schedules to efficiently respond to software and system issues that impact performance, deployment, and availability. These schedules ensure issues are not missed and are assigned to people with the skills to effectively evaluate and initiate a plan of corrective action.

Generally, a team of engineers is placed on "on-call" and if there are disruptions during a rotation, the team is notified. The best on-call schedules have redundancy with at least primary and secondary responders. If the primary responder is unable to take action in a designated time-frame, the secondary responder can be notified.

Modern on-call and alerting solutions like Opsgenie automate the notification process. Opsgenie uses flexible schedule options which are key for handling issues in a complex environment, and can be adapted from workflow to workflow. When there is an alert, on-call engineers get notified through channels such as SMS, voice call, push notification or email. As a notification is received, the alert is usually marked as acknowledged - meaning that someone got the notification and has started to work on this issue. Eventually, the on-call employees identify the problem and resolve it. As a best practice, after resolution, there is analysis to determine the cause of the issue and opportunities to avoid recurrence and accelerate response.



On-Call Concepts

WHAT IS AN ON-CALL SCHEDULE?

On-call schedules are used to determine who is on-call at a given time. On-call scheduling tools need to be flexible enough to support sophisticated scheduling scenarios as there can be very complicated scheduling requirements due to geographically distributed teams, schedule changes, team, and personnel changes.

In Opsgenie, an on-call schedule usually has one or more rotations and an additional rotation to support overrides. A **rotation** represents a set of people who rotate on-call responsibilities through the same shift. **Override** layers allow a user to replace an on-call user for a specified amount of time. **Final schedule** combines the rotations and override layers and shows the final schedule.

TODAY <	> No	ov 13 - M	Nov 26							1 Da	y 1 Wee	ek 2 Weeks	1 Month	Cale	ndar	Timeline
Rotations		_														
	11/1	3	11/14	11/15	11	/16	11/17	11/18	11/19	11/20	11/21	11/22	11/23	11/24	11/25	11/26
1 Weekday 👌	•	B	N	CFX	PP	BV	v	CFX			PP	BW CF	X PP	B	N	
1 Weekend a									SS	СК						SS CK
Overrides		ike on call	for an hour													
	11/1	3	11/14	11/15	11	/16	11/17	11/18	11/19	11/20	11/21	11/22	11/23	11/24	11/25	11/26
Final Schedule	9	⊷⊙														
	11/1	3	11/14	11/15	11	/16	11/17	11/18	11/19	11/20	11/21	11/22	11/23	11/24	11/25	11/26
Weekday's final		B	w	CFX	PP	BV	v	CFX			PP	BW CF	X PP	B	N	
Weekend's final									SS	СК						SS CK

In this on-call schedule example, we have five engineers: Bruce, Peter, Charles, Scott and Clark.

During the week, Bruce, Peter and Charles share on-call duty rotating daily, starting from 9:00 am until 9:00 am the next day. On the weekends, Scott and Clark are on-call, rotating daily from Saturday morning 9:00 am until Monday morning 9:00 am.

TODAY < >	Nov 13 - N	lov 26						1 Day	1 Week	2 Weeks	1 Month	Calen	dar	Timeline
Rotations +														
🖡 Daytime 🖋	11/13 A	11/14 A	11/15 A	11/16 A	11/17 A	11/18	11/19	11/20 A	11/21 A	11/22 A	11/23 A	11/24	11/25	11/26
1 Off-hours 🖋		ск	BW	RD	ск	вw			RD	СК	BW	RD	ск	
Overrides +	Take on call	for an hour												
	11/13	11/14	11/15	11/16	11/17	11/18	11/19	11/20	11/21	11/22	11/23	11/24	11/25	11/26
Final Schedule	⊚ ⊷													
Daytime's final	11/13	11/14 A	11/15 A	11/16 A	11/17 A	11/18	11/19	11/20 A	11/21 A	11/22 A	11/23 A	11/24 A	11/25	11/26
Off-hour's final		СК	BW	RD	СК	BW			RD	СК	BW	RD	ск	

In this example, the on-call schedule is based on business hours and off-hours. Here, "A" in the Daytime schedule represents the whole Application team. During business hours (week- days 09:00 - 18:00), all members of Application team are notified for alerts. In the evenings, Clark, Bruce and Raven share the on-call duty, rotating daily.

WHAT IS AN ESCALATION

Escalations determine who should be notified when there is an alert and what should happen if nobody responds to that alert. In other words, escalations are used to define who should be notified next, in which order, and through which timeline.



In this simple example, on-call engineers of the **DevOpsTeam_schedule** get notified immediately when a new alert is created. If this team does not respond to an alert in 5 minutes, second escalation step gets activated meaning that the next person on **DevOsTeam_schedule** schedule gets notified. If 10 minutes has passed since the creation of the alert and still no-one has taken any alert action, then the third escalation step is executed and all members of the **DevOpsTeam** team are notified.



The DevOps Movement and Its Crucial Use of On-Call Scheduling

Up until the last decade, responding to IT incidents was the primary job of the Operations teams. Organizations typically had a tiered team structure (i.e. Level 1, Level 2, Level 3) to respond to issues reported by customers or monitoring tools. The main goal in adopting this structure was to reduce operations cost. Usually, Level 1 would involve entry-level employees which was cost effective. If Level 1 was not able to resolve an issue, the issue would be escalated to Level 2, made up of more senior and therefore, expensive, people. This process would continue until the issue was resolved.

Recently, with the rise of "always on" services, customers expectations have risen. There are now greater interdependencies between systems. Urgency has become more important because a slow response to outages can impact a company's reputation due to social media. As a result of these factors, the structure of response teams needs to change and addressing incidents in a timely manner is mission critical. The DevOps movement set out to address this goal. Being most familiar with the code, developers are the ones who can best troubleshoot related issues in the shortest amount of time. Thanks to the DevOps movement, it is now common practice for developers to be on-call, ensuring the code runs well, lowering the MTTA and MTTR of alerts.

When the development team shares operational responsibilities, there are added benefits beyond the reduction in MTTR. Code is tested more rigorously, since that the team may be notified off hours if it has issues. In short, there is an ownership shift of mind and more reliable and resilient systems are built as a result.



Another reason on-call is becoming a more vital concept and on-call teams have to scale is the growth of organizations. As organizations grow, it is not just the number of employees that grow, but also the products, services, and the operational workload. Thus, it becomes a big challenge to continue to produce new features and grow products as well as make sure that the systems are healthy and up all the time. No surprise, all of these exciting complexities and changes naturally result in a change in organizations' on-call approach as well.

ONE OR MULTIPLE PEOPLE ON-CALL

In case of small organizations, usually, one person is on-call at a time and that person is responsible for the health of the whole application or system. As organizations add new people, expertise teams are usually built where each team is responsible for one or a few domains or microservices of the whole system. In the case of big organizations, there is a tendency to have on-call people from each team at the same time to make sure that there is good coverage in terms of the various skills and domains. Overall, this means having more people involved in on-call duties.

MULTIPLE ESCALATIONS

A very useful approach is to make sure there are multiple escalations available in each on-call schedule, so that the primary responders can get the help of others, as needed. This is especially good in case of growing companies, where there could be people on-call who are not experts of the systems yet. Here is an example on-call schedule with multiple escalations from our own usage of Opsgenie. We make sure that we have a backup schedule and that an experienced engineer is always available for additional support.

Integrations_Main_Schedule (GMT+03:00) Europe/Moscow MSK		₩ ℃ 4 0 0 ∨
TODAY < > Nov 13 - Nov 26	1 Day 1 Week 2 Weeks 1 Month	Calendar Timeline
Rotations +		
11/13 11/14 11/15 11/16 11/17 11/18 11/19 1 Main P PP JG CFX SS PP	11/20 11/21 11/22 11/23 11/24 JG CFX SS PP	JG
1 Integrations_Backup_Sched (GMT+03:00) Europe/Moscow MSK		₩ ℃ 4 1 ∨
TODAY < > Nov 13 - Nov 26	1 Day 1 Week 2 Weeks 1 Month	Calendar Timeline
Rotations +		
I Backup A BW CK BW CK BW CK BW	11/20 11/21 11/22 11/23 11/24 CK BW CK BW	4 11/25 11/26 CK

ESCALATING TO A DIFFERENT TEAM

Another pattern we observe from our customers' usage is related to escalations. There is usually a first escalation of unattended issues to the back-up people in the schedule, then to other people in the same team. If the issue is still not taken care or resolved, another team helps. In the example below, unresolved issues of **Database** team get escalated to the **Platform** team.

‡ Database_Team_Escalation			ළු	Gant	Û	~
immediately	🛗 On call users in Databa	ase_Team_Schedule, If not acknowledged				
5 min.	mext user i	in Database_Team_Schedule, If not acknowledged				
10 min.		P All members of DBA, If not acknowledged				
	20 min.	👂 Routed target of Platform, If not acknowledged				

SEPARATE NIGHTWATCHERS

One pattern we see in our customers' usage is to have different set of people on-call during business hours and off-hours.

1 Database_Tea	m_OnCall_S	Giche (G	MT+03:00) Euro	pe/Moscow M	SK								# 2 4] / 🛍 🗸
TODAY < >	Nov 13 - M	ACCAIL_Sche (GMT+03:00) Europe/Moscow MSK xv 13 - Nov 26 a 11/14 RD RD RD 11/15 11/15 11/16 11/18 RD						1 Day	Calen	dar	Timeline			
Rotations +	۱. ₁ .													
1 Daytime 🖋	11/13 RD	11/14 RD	11/15 RD	11/16 RD	11/17 RD	11/18	11/19	11/20 NR	11/21 NR	11/22 NR	11/23 NR	11/24	11/25	11/26
1 Off-hours 🖋		ск	СК	CK.	ск	Clar	k Kent		BW E	BW E	BW E	BW	Bruce W	ayne
Overrides +	Take on call	for an hour												
	11/13	11/14	11/15	11/16	11/17	11/18	11/19	11/20	11/21	11/22	11/23	11/24	11/25	11/26
Final Schedule	⊚ ⊷⊙													
Daytime's final	11/13	11/14 RD	11/15 RD	11/16 RD	11/17 RD	11/18	11/19	11/20 NR	11/21 NR	11/22 NR	11/23 NR	11/24 NR	11/25	11/26
Off-hour's final		ск	ск	СК	ск	Clar	k Kent	E	3W I	BW E	BW I	3W	Bruce W	ayne

In the example above, note that the nightwatcher is on-call during the weekends as well.

SCHEDULES WITH OVERLAPS

Some companies make sure there are overlaps between shift changes so that the previous on-call engineers can have some time with the next ones and can delegate the work and communicate properly.



FOLLOW THE SUN SCHEDULES FOR INTERNATIONAL ORGANIZATIONS

Some companies have international teams and can have 24/7 coverage easily within business hours, as teams have different time zones. In the example below, there are on-call engineers both located in Europe, in the U.S. and in Australia. An engineer from each location covers the on-call duty during business hours, each for 8 hours a day.

1 On_Call	_With_	Handover	(GMT+03:0	00) Europe/Mo	oscow MSK									†		Ìv
TODAY	< >	Sep 25 -	Oct 8						1 Day	1 Week	2 Weeks	1 Month	Calen	dar	Timeline	
Rotations	+		.													
1 Rot1	ø	9/25 BW	9/26 BW	9/27 BW	9/28 BW	9/29 BW	9/30	10/1	10/2 CK	10/3 CK	10/4 CK	10/5 CK	10/6 CK	10/7	10/8	
1 Rot2	ø		RD	RD	RD	RD	RD			CFX	CFX	CFX	CFX	CFX		
1 Rot3	A ¹						S	cott Summers							James Howlet	tt

In the example below, there are on-call engineers in India and in the U.S. Each engineer has on-call duty for 12 hours a day.



Advice to Management Teams While Enrolling Changes to On-Call Systems

Being on-call can be a daunting and disruptive experience. Many people with on-call duties complain how having to be ready to handle incidents affects work life balance, even health, as on-call employees may be frequently woken up in the middle of night or may need to plan evenings and weekends while considering on-call duties. As organizations enroll changes to scale on-call teams, it needs to be considered how to best match that evolution with a sustainable and humane solution. Below is some advice based on our experiences so far with our customers.

1 TRANSPARENCY

Transparency is the key to successful communications. Clarifying expectations around availability of employees is a must when rolling out an on-call system or a change to an existing on-call system. People with new on-call duties may have questions like the following; the answers should be clear before enrolling the changes:

- Are engineers supposed to be on-call during nights?
- If on-call during nights, is there flexibility to work from home the next day?
- Or start the next day later than usual to make up needed sleep from the on-call part of the night?
- Are engineers supposed to do development work during on-call time?
- Maximum how many times in a month would an engineer be on-call?



TRAINING

Providing proper training to new on-call engineers is a must for scaling on-call systems successfully. Having an established training plan for on-call processes and tooling is the first step in this direction.

Providing engineers with detailed and up-to-date runbooks is also very important. That helps engineers see what the protocol is for certain situations, or how similar issues were resolved effectively in the past.

Shadowing experienced on-call engineers would be a valuable experience as well, so that the new on-call engineers can feel the on-call atmosphere directly while still being given some guidance and support.

Having multiple escalation channels on on-call schedules is a good idea especially for junior on-call engineers. Not being completely alone when help is needed will come as a relief. Related to this, having the junior engineers in the primary on-call rotation and having a more senior engineer as back-up or in a secondary rotation is another best practice. This organization helps junior engineers develop the required on-call skills while avoiding panic when there is an issue beyond junior expertise.

3 DEVELOPMENT DUTIES DURING ON-CALL

Having development duties during on-call usually means lots of context switching and interruptions, especially when the on-call duty is a heavy one. It also has negative effect on development sprints since it is difficult to estimate how much contribution on-call people may have to the development sprint. This usually means less development efficiency and more stress on the on-call engineers. As a best practice, we recommend not assigning development duties to developers during on-call duties and when there is free time, request for the developer to work on improving the on-call related documentation and automation to eventually improve the sustainability of the systems and services.

WELL-DEFINED AND FINE-TUNED PROCESSES AND SYSTEMS

A healthy on-call system can only exist if improved constantly by fine-tuning processes and systems. A few recommendations towards this goal:

- Evaluate alert priority and urgency and set systems based on that.
 Low urgency alerts can wait until the morning; on-call employees do not need to wake up in the middle of the night because of these alerts.
- Reduce false-positives by classifying alerts based on factors such as root-cause, originat ing system, message, thresholds etc. This way, differentiate actionable alerts from the rest.

- Deduplicate related alerts to avoid alert fatigue; on-call employees do not need to get notified by a flood of alerts in an hour for the same issue; give the chance to focus on the resolution instead of notifying with the same alert over and over.
- Design rich alerts that empower the on-call engineers to make effective decisions and apply the knowledge recorded in runbooks.
- Provide alert reports and metrics to on-call teams so that the weak areas can be seen in the systems and insights gained on the root-causes and resolutions. Do not let on-call teams get bogged down by the same problem repeatedly.

5 REVIEWING ON-CALL REPORTS

For fairness, managers are advised to review on-call related reporting and check how often each team member was paged or woken up, and enroll changes in the on-call system considering such facts and hopefully avoid employee burn-out. Below is an on-call report for one of our teams internally. The report shows information such as how long each team member has had on-call duty, distribution of on-call hours for each on-call schedule, and the hourly and daily distributions of oncall duty for each person.







Platform Non-Work Hours Schedule
 Platform Deployment Schedule
 Platform Non-Work Hours Backup Schedule
 Platform Work Hours Schedule
 Platform Work Hours Backup Schedule

6 FRIENDLY ON-CALL CULTURE

On-call engineers carry huge responsibility for the success of companies. The job requires extra caution when rolling out changes, and responding knowledgeably as fast as possible when there is a problem. Such responsibility may naturally mean stress and tension, especially when there is a big issue and there are unknowns. The on-call culture set by the senior on-call engineers and management teams defines how people deal with that stress and tension and whether the on-call experience is transformed for the good ofthe company or not.

Lastly, we recommend management teams make sure that on-call employees are being listened to. Management should organize all-hands meetings with the on-call engineers to discuss problems, complaints, and areas of weaknesses. Make sure to take actions which help resolve these problems and weaknesses. Constantly reevaluate the on-call system, tools, processes, people, documentation, and training plans that are being used. Lastly, make sure that management is not the sole decision maker when it comes to on-call organization and protocol, and instead, make each decision a team decision.

For both the on-call engineers' sake and the on-call culture of the company, management teams should pay attention to developing a friendly on-call culture and make it clear that the goal should always be to find the problems, risks, and weaknesses in systems and solve them. The goal should never be to find whom to blame. Opsgenie offers the tools to make on-call schedules, processes, and experiences easier and provides all the opportunities to continue improving at every turn.

