

Cadence

Increase Efficiency by Using a Rhythm

By Kathy Iberle

Key Points

- **A cadence is a regular rhythm within a process.**
- **Putting a recurring activity on a cadence reduces overhead.**
- **The best cadence for an activity is a trade-off between the value of doing the activity more frequently and the overhead incurred by doing the activity.**

What is a Cadence and Why Should I Have One?

A cadence is a regular, predictable rhythm within a process. For instance, a particular meeting happens on the same day of every week, or your website is refreshed on the last day of the month every month.

A cadence saves time by reducing overhead. Instead of coordinating multiple people's availability and finding a meeting room each time a meeting is held, the calendar coordination and meeting room reservation is done just once.

A cadence also saves time by reducing waste. When an activity happens at a predictable time, people can plan their work to take advantage of the cadence. When a bus runs on time reliably, you'll go to the bus stop just in time to catch the bus. If the bus arrives erratically, you'll go to the stop early – and probably waste some time waiting at the bus stop.

Which Activities Might Be More Efficient on a Cadence?

Some activities which typically benefit from a cadence are:

- planning and scheduling
- integration across multiple teams
- work by anyone who is assigned to your team 10% or 20% time
- progress updates and workload balancing
- product releases

When coordinating or planning an activity requires a lot of back-and-forth conversation among many people, you may all be better off to put that activity on a cadence.

For instance, imagine your team is allotted 10% of an expert's time. Everyone on your team sends questions to the expert at random times during the week, and he calls your people back or arranges quick meetings with them to answer their questions. The questions are answered quickly, but the expert spends a quite substantial portion of your precious 4 hours each week on leaving messages and arranging times to meet, not to mention the interruptions as the questions arrive in the midst of his other work.

You'll get much more value from the expert if you put his work on a cadence. Try scheduling two regular time-slots per week for the expert to visit your team, and ask your team to bring all their questions to those time-slots. This reduces the overhead for the expert, allowing him to spend more time on your issues. There's a bonus, too – the expert and your team have more opportunities to see relationships between the issues which might be missed otherwise.

How Do I Create a Cadence?

Creating a cadence can be as simple as setting up a recurring meeting to do specific work at the same time every week.

The first step is identifying the activity (or activities) which will be put on a cadence. The activity must be something concrete that has a clear output or deliverable, such as questions answered or a decision made.

Then schedule a recurring date on a regular cadence, such as once a week or every 4th Tuesday. Set up whatever guidelines or rules are needed to discourage breaking the cadence. If the cadence is not reliable, people will keep doing work outside the cadence and not get the benefit of the cadence.

Finally, consider whether there are related activities which should also be put on the same cadence or a coordinated cadence. For instance, a sprint in Scrum puts several activities on the same cadence – integrating user stories together for delivery, demonstrating user stories to the customer, and deciding which user stories to start next. All three activities are done once per sprint.

How to Determine the Best Cadence for an Activity

The ideal cadence for an activity is a balance between:

- The cost of doing the activity (known as *transaction cost*)
- The cost of not doing the activity (known as *holding cost*)

Most people pay attention to the transaction cost and vastly underestimate the holding cost. One of the early arguments against iterative methods such as Scrum was that it would be expensive to run regression tests on each sprint, rather than just once at the end of the project. That's a *transaction cost* - it gets paid every sprint, and it does add up.

However, the *holding cost* of **not** running regression tests early, and **not** finding the bugs early, is usually quite a lot higher. The holding cost is much more difficult to see, however.

To determine a good length for your cadence:

1. Make a list of the valuable output from the activity. **Do** include information as output – information is often the main output of many activities. For instance, one output of regression testing is information about defects.
2. Consider how often you would be able to use this output. If you could get it weekly, would that improve your ability to avoid unneeded work? How about daily? How about every hour? A cadence can be too fast!
3. Consider how long it takes to create a new piece of output. You may not be able to get new information daily.

Now, choose a cadence that is short enough to get valuable output as often as you can realistically use that output. Try it and see if it works for you.

Cadence – The Bottom Line

Cadence is a simple project management tool that is often under-utilized. Many activities are less expensive when put on a regular rhythm or cadence. Shortening the cadence can have surprisingly good results.

By the way, there is a formal way to analyze the trade-off between transaction cost and holding cost. You can see the math in all its glory in my white paper [“Transaction Cost, Holding Cost, and Cadence: Finding the Optimum Frequency for a Recurring Activity”](#). Or, read the sections on “Cadence” and “Cadence in Action” in *The Principles of Product Development Flow: Second Generation Lean Product Development*; Reinertsen, Donald G.; 2009.