

Draft Proposal for TSC Election Process

Click for anchor link [Shall the TSC approve the updated proposal for the 2021 Election Mechanic as written below](#) This vote macro is locked

Choices	Your Vote	Current Result: (7 Total Votes)
+1	<input type="checkbox"/>	<div style="width: 100%;"><div style="width: 100%; height: 10px; background-color: #00a0c0;"></div><p>7 Votes , 100%</p></div>
0	<input type="checkbox"/>	<div style="width: 0%;"><div style="width: 0%; height: 10px; background-color: #ccc;"></div><p>0 Votes , 0%</p></div>
-1	<input type="checkbox"/>	<div style="width: 0%;"><div style="width: 0%; height: 10px; background-color: #ccc;"></div><p>0 Votes , 0%</p></div>

** Thanh Ha (zxiiro)voted for Tejas Nevrekaras Proxy.

- Click for anchor link [Shall the TSC approve the updated proposal for the 2021 Election Mechanic as written below](#) This vote macro is locked
 - ** Thanh Ha (zxiiro)voted for Tejas Nevrekaras Proxy.

Overview

Principles

- [Representation of Constituencies](#)
- [Election by Community Peers](#)
- [Protection from Dominance by a Constituency](#)
- [Fairness of Exclusion](#)
- [Modern Election Tools](#)
- [Flexibility over Time](#)

Frequency

Membership Definitions

Seat Allocation

Election Mechanics

- [Election Algorithms](#)
- [Over-Max Resolution Through Scaled Popularity](#)
- [Scaled Popularity Example](#)
- [Represented Groups](#)
 - [TSC Seat Allocations By Represented Group for the 2021 elections](#)

Overview

This page is the result of the community's effort to refactor the OpenDaylight Technical Steering Committee (TSC) election system. It provides background information, guiding principles, a "framework" of ideas that have consensus. OpenDaylight's Technical Steering Committee will be elected by a set of multi-winner elections, one election per Represented Group. The set of Represented Groups and the number of TSC seats allocated to each will be determined before each election by the outgoing TSC.

Principles

This section documents *why* we designed the election system we did.

Representation of Constituencies

OpenDaylight is made up of groups that sometimes want different things from the TSC. These Constituencies should be represented at the TSC.

Election by Community Peers

Representatives to the TSC should be elected by their peers in the OpenDaylight community on the basis of merit.

Protection from Dominance by a Constituency

No single Constituency should be able to dominate the TSC.

Fairness of Exclusion

Candidates should be excluded from the possibility of holding a TSC seat they would have otherwise won as infrequently as possible, and only for the sake of trade-offs with other key principles. When some candidates must be excluded from holding seats, the selection should be done as fairly as possible (consensus, merit, random, not power over career).

Modern Election Tools

When possible, OpenDaylight's election process should make use of high-quality existing election primitives. For example, proportional representation voting systems for multi-seat constituencies that use ranked voting, like [Condorcet](#) or [Single Transferable Vote](#), have well-understood good properties that are helpful for our use-case.

Flexibility over Time

We have an imperfect understanding of what governance structures are best for OpenDaylight and we expect the governance needs of OpenDaylight to change over time. The election system we adopt should enable a clear path for adjustments over time, especially for aspects that are most likely to be dynamic (like seat allocations to Constituencies).

Frequency

TSC elections will happen [yearly](#).

Membership Definitions

- Active Community Members: Anyone from the OpenDaylight community with twenty (20) or more measurable contributions during the previous 12-month period, inclusive of code merged, code reviews performed, wiki page edits, or JIRA activities.

Seat Allocation

OpenDaylight TSC seats are allocated to Represented Groups (RGs) of the OpenDaylight community. Represented Groups are typically Constituencies with systematically competing interests that the TSC wants to ensure have voices (in the form of votes). Seats may be allocated to other groups, for example, to incentivize behaviors like projects moving to later [ODL Lifecycle states](#).

Before each election, the outgoing TSC will review and potentially make changes to the allocation of TSC seats.

The seat allocation for an election is defined as a set of Represented Groups, each of which is defined by:

1. Min Seats - Minimum number seats allocated to the RG (as an integer or ratio of total). Provides Representation of Constituencies property.
2. Max Seats - Maximum number seats that can be held by the RG (as an integer or ratio of total). Provides Protection from Dominance by a Constituency property.
3. Candidates - Set of people eligible to run for the RG's seats. Provides Election by Community Peers.
4. Voters - Set of people who can vote in the RG's election. Provides Election by Community Peers.
5. Duplicate Voter Strategy - If a person meets the Voter description of this RG for more than one reason, do they get a single vote or a vote for every interest they represent? Resolves Multiple Votes in Same Represented Group corner case ([TSC vote](#))

Election Mechanics

OpenDaylight will run a series of elections, one per Represented Group. Each election will have N winners, where N is the Represented Group's "Min Seats" value.

For each Represented Group, the eligible candidates for its seats should be notified of their eligibility and asked to self-nominate. The TSC or some Election Official appointed by the TSC will specify a location to post candidate information (like a wiki), and may additionally suggest (biographical, platform) or require (name, IRC handle) some discretionary information. Self-nominations from candidates must include the set of Represented Groups for which they match the Candidate description and as well as a single Represented Group under which they would like to run for a TSC seat.

The self-nomination phase should remain open for a period that enables eligible candidates from all Represented Groups to complete the process with a reasonable effort. If a Represented Group fails to produce enough self-nominated eligible candidates to fill the seats allocated to it, the TSC should reopen the self-nomination phase, encourage additional nominations and potentially revise the seats allocations for the election.

After the self-nomination phase, for each Represented Group, the eligible voters in its election should be directed to the candidate information and given a ballot for each Represented Group for which they are a member of the electorate. If the Ranked Preference Algorithm tooling supports it, ballots will list all self-nominated eligible candidates in a pseudorandom order. Exact wording for the ballot and other communication can be specified by the TSC or left to an Election Official. For RGs with a Vote-per-Interest Duplicate Voter Strategy, voters are given a ballot for each reason they match the RG's Voters description (may require additional email addresses). For RGs with a Vote-per-Person Duplicate Voter Strategy, no voter receives more than one ballot. Voters vote and the Ranked Preference Algorithm returns the ranked preference results. For each RG, the top "Min Seats" candidates are declared winners pending over-max resolution.

For all Represented Groups, the number of TSC seat winners who meet the Candidates description of the RG (declared during self-nomination) is determined. The actual percentage of representation achieved by each RG is determined. If the actual percentage of TSC seats won by candidates in a Represented Group is higher than the group's Max seat allocation, over-max resolution is handled by a Scaled Popularity Election.

Winning candidates are notified, results are publicly posted and the transfer of TSC seats happens immediately.

Election Algorithms

OpenDaylight's Election Framework allows the specific ranked preference voting algorithm it consumes to be swapped out. This gives the TSC or their Election Official flexibility over time, allowing the best-known algorithm and tooling to be used for each election. This also allows the algorithm and implementation to be treated as a black box in the rest of the Election Framework, simplifying user-facing elements by extracting complexity to modern election tools.

For the sake of the rest of the Framework, the specific algorithm and tooling used for a given election will be called the Ranked Preference Algorithm. The properties of modern election algorithms are complex, so the TSC and their Election Official should generally strive to follow modern best practices. The only specific requirement imposed by the way the algorithm is consumed in OpenDaylight's Election Framework is that it must provide a ranked preference result, so the top N candidates can win the N seats in their Represented Group's election.

The [Condorcet algorithm](#), with the Condorcet-IRV completion rule, will be used for OpenDaylight's first new-style election. Tooling will be provided by the [Condorcet Internet Voting Service](#), which is [open source](#) and maintained by Cornell. ODL's Elections Framework was originally designed with STV in mind, but no suitable web-based implementation can currently be found and Condorcet may be a more fair algorithm.

Over-Max Resolution Through Scaled Popularity

The TSC [voted](#) to use Scaled Popularity for over-max resolution.

After running the elections for all Represented Groups, for any RG that exceeded the Max seat value, an over-max resolution election is held (using the votes from the full election, no new voting needs to occur) among the winning candidates. Each candidate's vote count is scaled to account for differences in the number of votes cast in their Represented Groups. The resulting counts are used to run a normal Ranked Preference Algorithm election among the over-max RG's winning candidates, with the number of winners equal to the RG's Max seat value. After all losing candidates have been determined, they are dropped from the main Represented Group elections they competed in. Elections for all Represented Groups are then re-run. If the new results put an RG above its Max Seats cap, the process repeats until all RG are below their cap. If the pool of self-nominated eligible candidates is exhausted before the process terminates, the outgoing TSC should revisit the company cap size or recruit a more diverse set of candidates.

Scaled Popularity Example

As an example, if we had three elections with 4, 8, and 16 voters in each, where we wound up with four candidates from RG A winning, but a cap of three, we would have a special side election to figure out which one should be dropped. This would be done by scaling the candidate's votes by $1/n$, where n is the number of voters in the RG, and re-running the Ranked Preference Algorithm election.

So, if the candidates got:

- A.) 2 votes of 4
- B.) 3 votes of 8
- C.) 5 votes of 16
- D.) 3 votes of 16

That would result in candidates having scaled votes of:

- A.) 0.5
- B.) 0.375
- C.) 0.3125
- D.) 0.1875

The result is that person D is eliminated and the main RG elections are re-run.

Represented Groups

TSC Seat Allocations By Represented Group for the 2021 elections

The TSC selected the "[Only Committer-At-Large \(CAL\)](#)" set of RGs. The TSC also decided that [at most 49% of TSC seats be given to employees of a single company and it's affiliated companies](#), but did not cap other RGs.

- Committers-at-Large
 - Min Seats: 5
 - Max Seats: 5

- Voters: Committers to ODL projects
- Duplicate Voter Strategy: Vote-per-Person
- Active Community Members
 - Max Seats: TBD
 - Voters: Active Community Members
 - Duplicate Voter Strategy: Vote-per-Person

For purposes of avoiding dominance from a particular company, also define any given company as a Represented Group with no seats (and thus no need to define voters):

- for company_var in <the set of all companies>:
- Represented Group company_var * Min Seats: 0 * Max Seats: 49% (cap of 5)
- Candidates: Works for company_var

Total seats: TBD