Backbone Change Management

Scaling @ Million+ Operations a Year

Karthik Subraveti

Meta

Why is this difficult?

Backbone networks are **complex** and **diverse**, with several constraints. Performing operations is typically considered **high risk**.

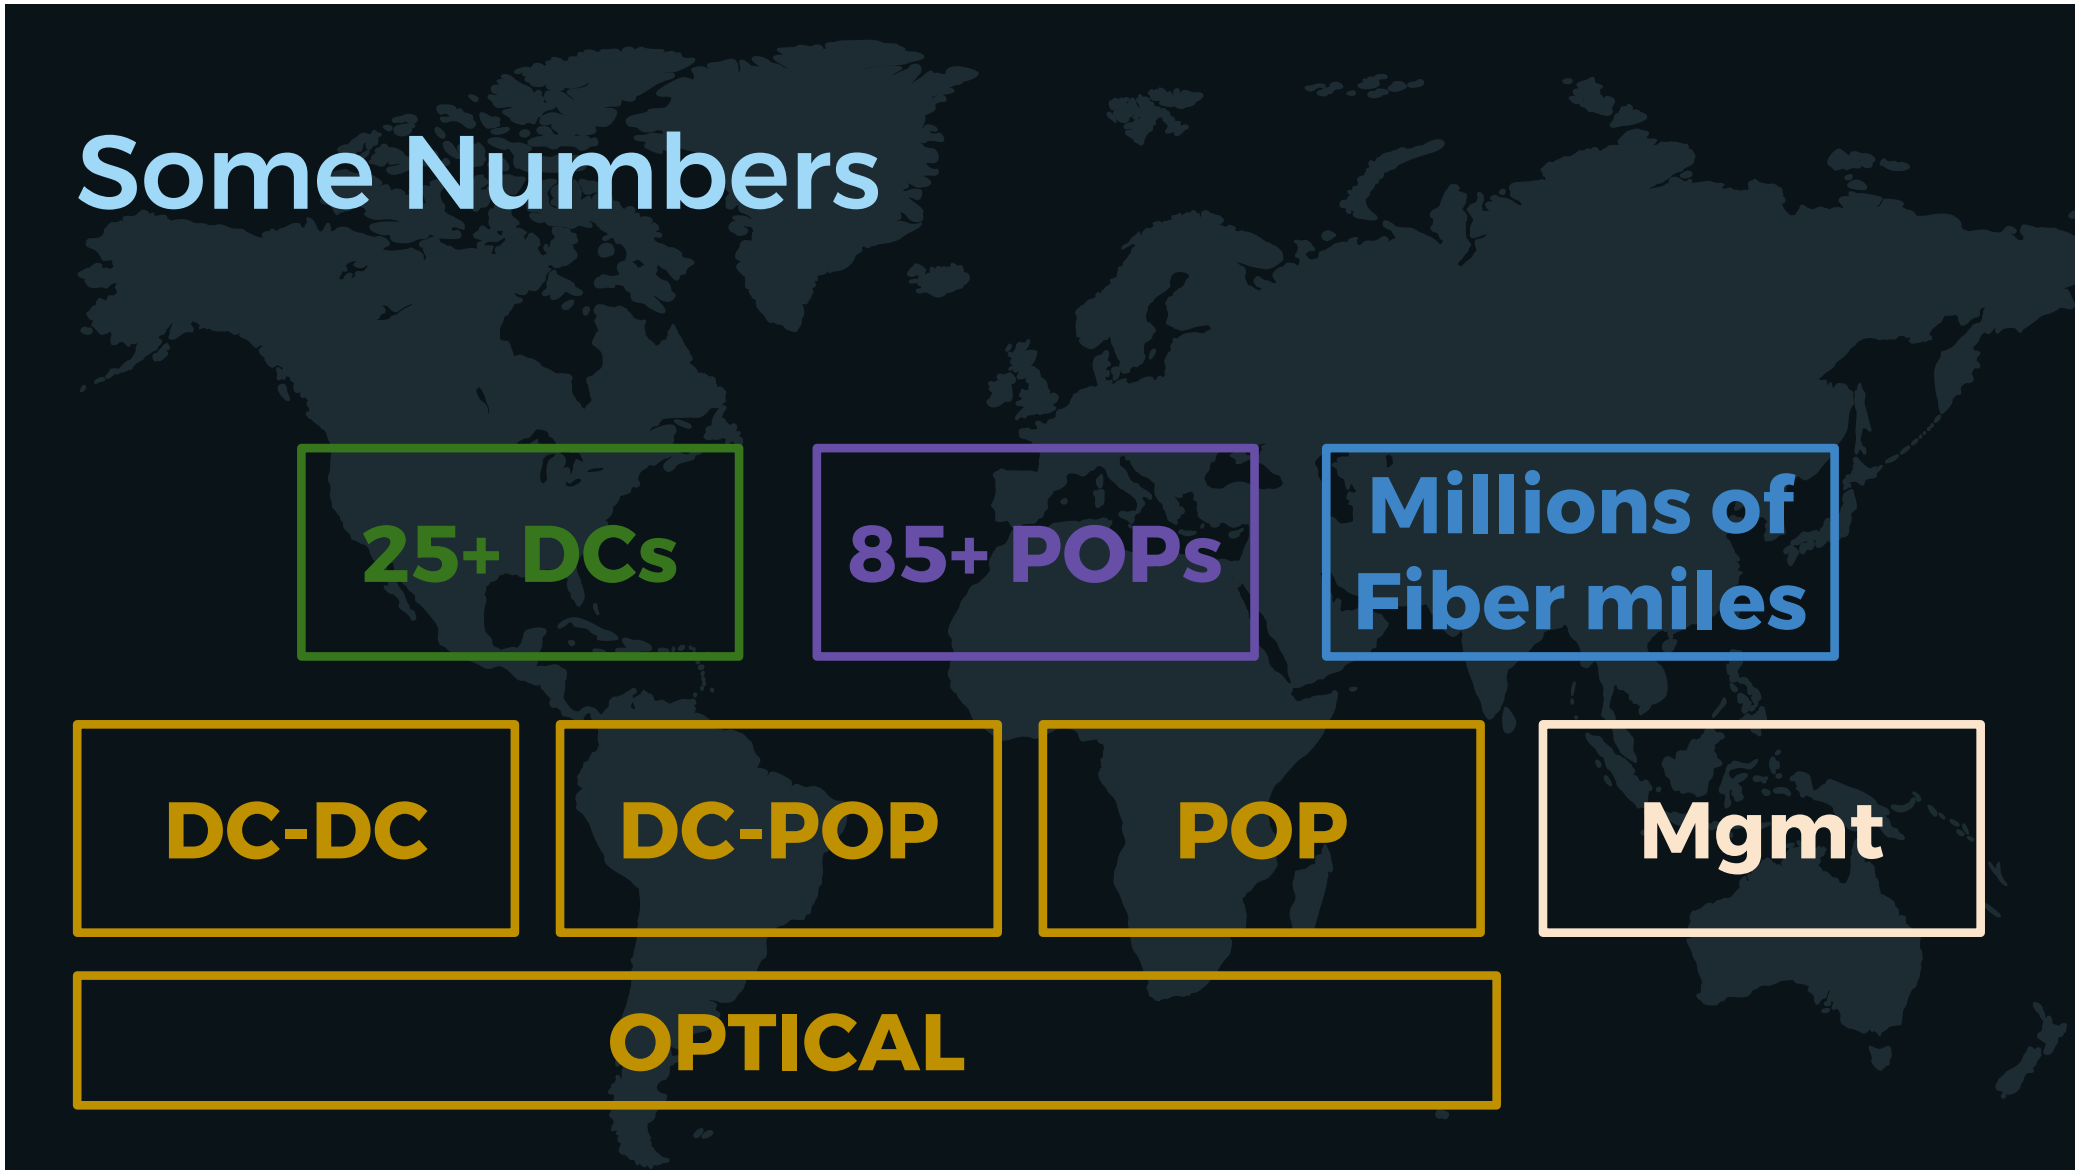
Some Numbers
25+ DCs
85+ POPs
Millions of Fiber miles
DC-DC
DC-POP
POP
Mgmt
OPTICAL

Story time - Long long ago...

...operations were less complex, teams were small. Work was easily coordinated.

As networks grew, so did operational complexity and teams.

#1 MVP

Goals

1. **Central** place for viewing operations
2. Block **risky** operations
3. Support **external** vendor maintenances

Anatomy of an Operation

Network Entity

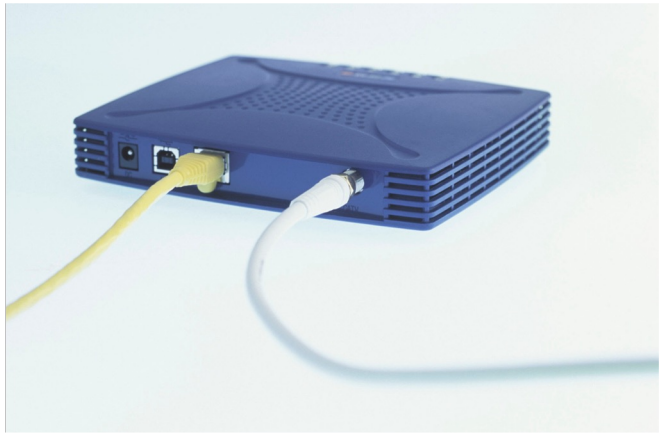

Start / End

Operation Type

Scheduler - *Introduction*

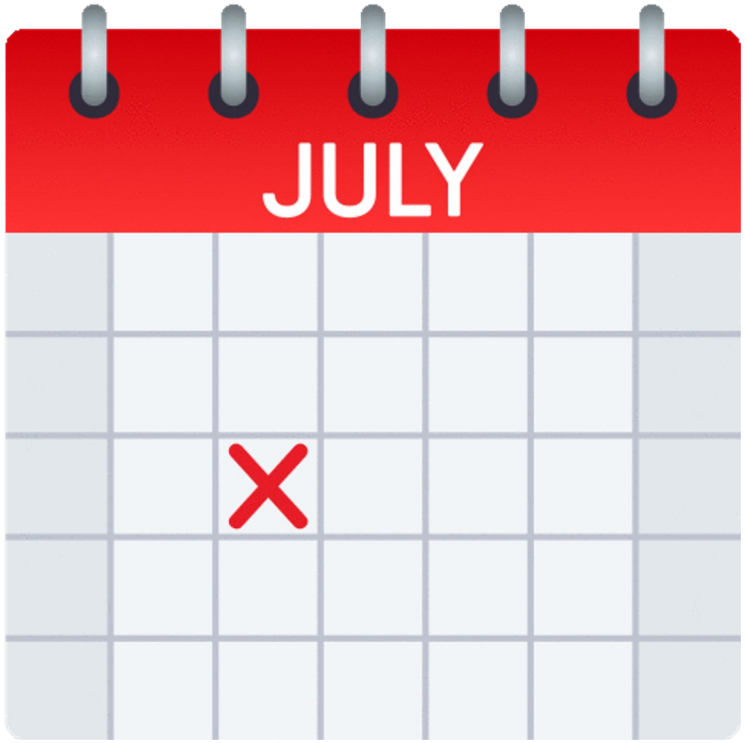

- Central place for storing all operations
- Schedule operations based on risk
- Automatically trigger traffic shift(drain) prior to operation start

Scheduler - *Working*

- Reads active operations and outages, assessing risk.
- Resolves risk by prioritizing runnable operations and blocking others.

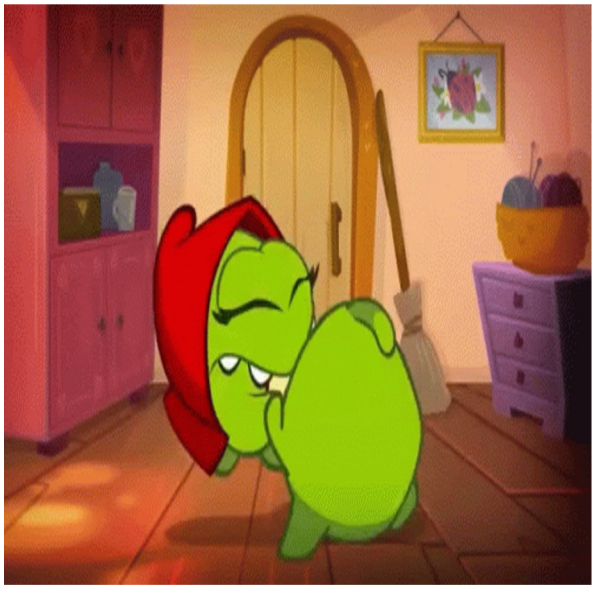

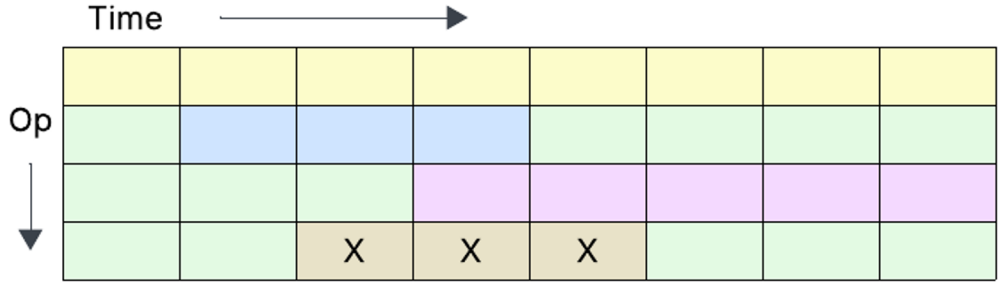

Scheduler - *Scope Checks*

- Scope checks are limits on number of operations performed within device groups.
- This is critical to containing the blast radius of operations.

Scheduler - *Risk Checks*

Verifies network **capacity** exceeds **demand** through simulations, including **Maxflow** and **Protected Capacity**, to ensure safe operations and minimize risk.

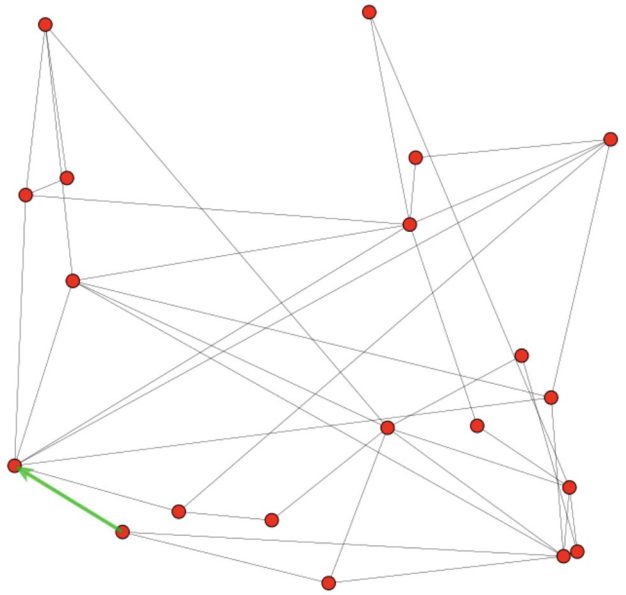

Story time - Reconfigure the World

Perennial problem of never reconfiguring the devices resulting in 100s of devices with stale configs and images

#2 Fully Automated Fleet-Wide Operations

Orchestrator

- One device batch per operation.
- Manages operations across several device batches.
- Support health checks

Story time - Gold Traffic at Risk

- Concurrent fiber maintenances on 2 out of 3 paths
- Backbone network suffers significant loss.

And more.....

#3 Build Early Warning System

Evaluator

- Establish operational policies for safety checks.
- Identify potential risks 2 weeks in advance and de-escalate

Story Time - The Missing Operations

- Large scale migrations delayed by unexpected operations
- Decomm accidentally performed on live circuits.

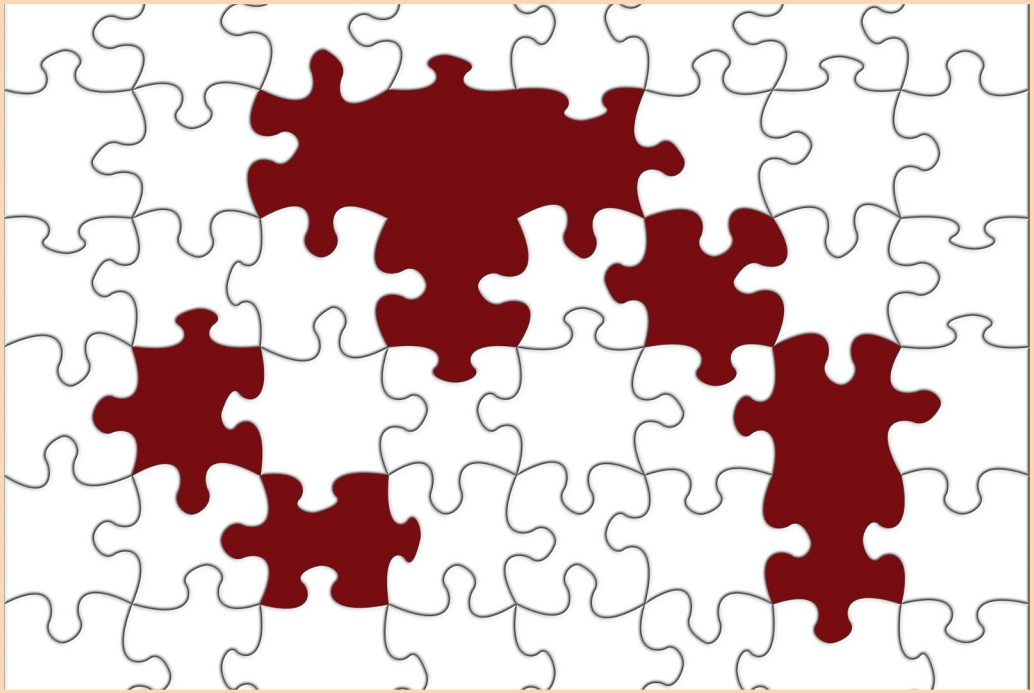

#4 Onboard all Operations

Intent Aware Scheduling

- Treating all operations equally no longer works.
- Planned operation may require relaxed scope checks, while fleet-wide upgrades need stricter controls.

Story Time - Noisy Health Checks

- Introduced to ensure smooth operations, health checks became noisy and intrusive. Unrelated failures (e.g., ACL push blocked by line card issues) caused frustration.
- Lack of clarity on where to add checks led to duplication and confusion

#5 Evolving Health Checks

Workflow Specific Health Checks

- Support added to include workflow specific health checks in addition to baseline checks.
- Workflow specific health checks added in orchestrator to pre-operation, post-operation, and after the device is returned to production.

Putting it all together

Orchestrator

Workflow specific health checks

Orchestration of change workflows

Scheduler

Intent aware scheduling

Evaluator + Operational Policy

Safely schedule operations, auto shift traffic

Impact

- More than a million network operations on a yearly basis are done safely without major incidents
- Less than 0.5% human involvement
- Up-to-date fleet with latest image and configs
- Optimized repair resulting in lower MTTR

This Journey is 1% Complete

- Simpler abstractions don't work for large-scale migrations. Need to support newer abstractions.
- Even 0.5% human involvement is too high, need newer ways to reduce complexity.
- Explainability is a major challenge.

Takeaways

- Backbone change management is iterative, improved through continuous work and partner feedback.
- Metrics are essential for identifying system gaps and driving improvements.

Questions?