


OPEN

Compute Project





OCP U.S. SUMMIT 2016

March 9-10 | San Jose, CA

Platinum Sponsor

OCP U.S. SUMMIT 2016

Facebook switch software development

Jasmeet Bagga
SOFTWARE ENGINEER


Outline

- ➔ Overview of wedge (FB ToR switch)
- ➔ Software evolution from ToR to modular switch (6-Pack)
- ➔ FB software development methodology

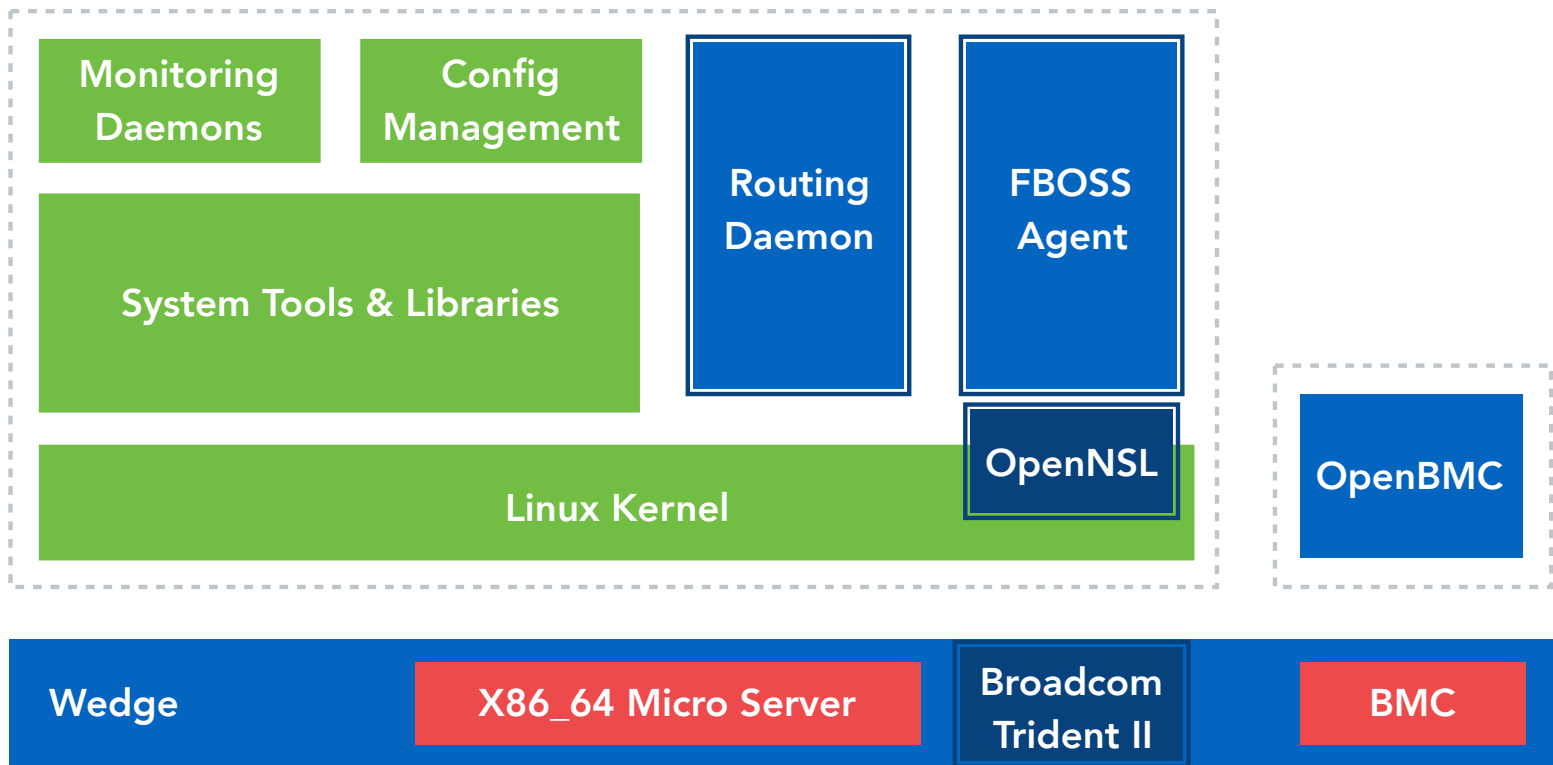
Let's build a TOR (2013)

- ➔ Satisfy a small feature set and do it well
- ➔ Move fast on desired features and bug fixes
- ➔ Extensible API via thrift to integrate with custom controller applications
- ➔ Build it like a server to leverage existing components


FB Server



Wedge



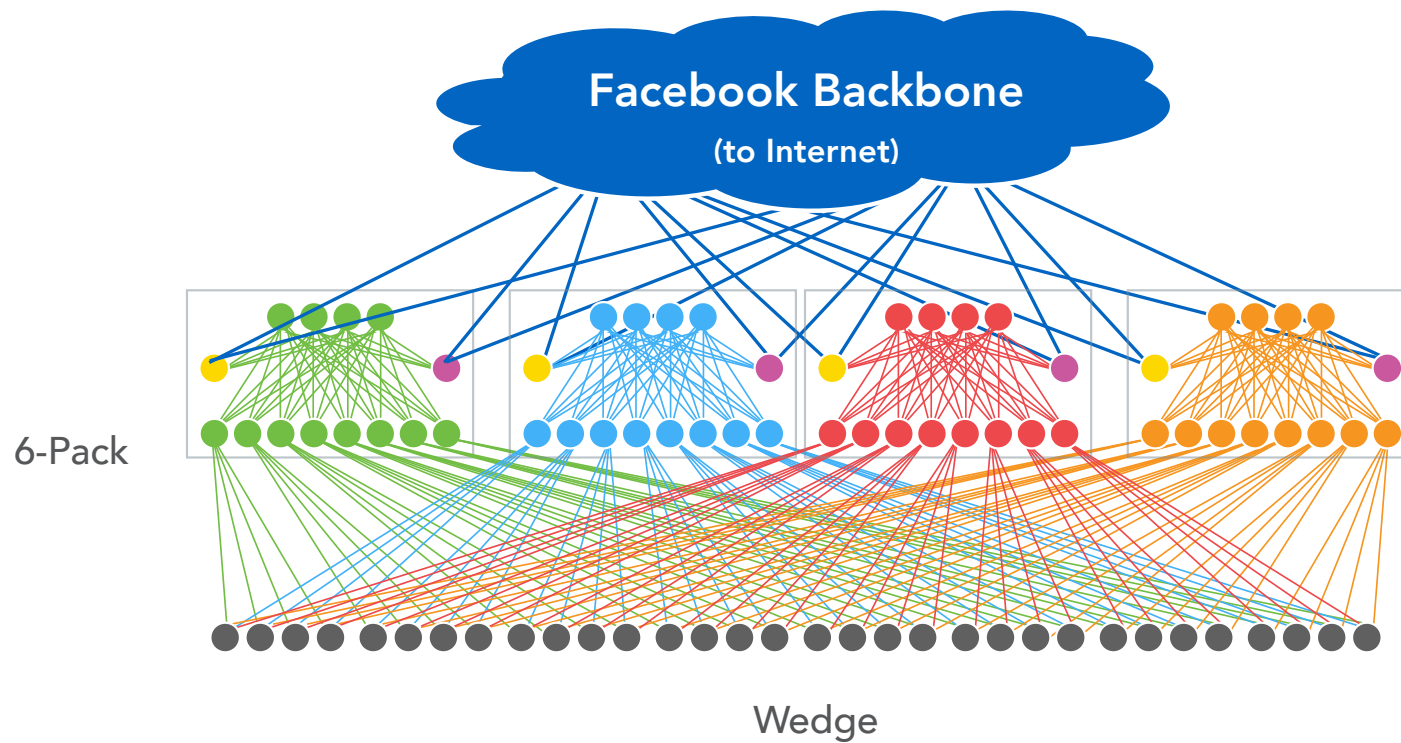
Success !!



From Wedge to 6-Pack



Moving up layers



A large, fiery nuclear explosion with a massive mushroom cloud rising from a dark, smoky landscape. The sky is filled with dark, swirling clouds, and the ground is covered in a layer of ash or smoke.

Complexity scale & explode


Routes: 1K vs 16K

Failure domain: 1 rack vs 48 racks

Traffic: 640Gbps vs 5.12Tbps

BGP peerings: 4 vs 112



No central brain



Divide & conquer

→ Just a set of wedges connected together

→ Outside \sim Inside



Climbing mountain 6-Pack,

- ➔ Fitting 16K routes on ASIC
- ➔ Route programming performance
- ➔ ECMP balancing
- ➔ Failure tolerance





Routes everywhere

We are out of TCAM

A close-up photograph of a snail moving across a dark, textured gravel surface. The snail's shell is brown and black with a spiral pattern. Its body is a pale, translucent yellowish-brown. Two long eye stalks are extended forward. The background is a soft-focus, light greenish-grey.

Hey slowpoke

300s => ~5s to program/delete 16k routes

A photograph of a yellow fire hydrant on a grassy area next to a road. Water is leaking from the side of the hydrant, creating a large puddle on the grass and road. The scene is used as a metaphor for resource waste in a technical context.


So wasteful

ECMP using only half the links when
wedge and 6 pack are connected

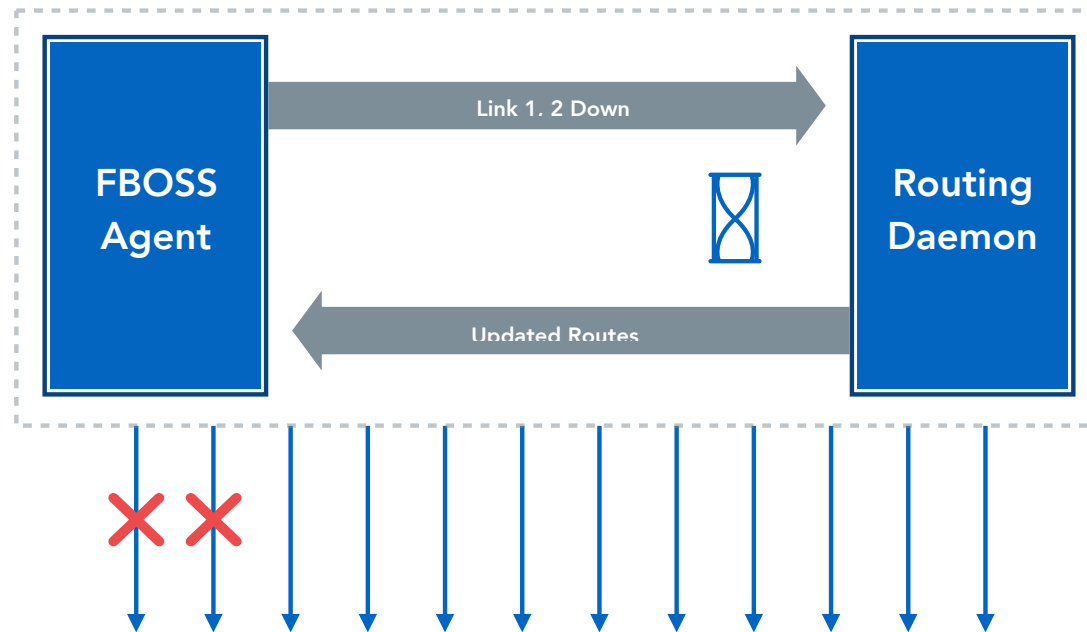
A photograph of a forest scene where a large, thick tree trunk has fallen across a narrow dirt path. The trunk is dark and textured, with its branches extending into the background. The path is muddy and wet, with some green moss or algae visible. The surrounding area is filled with lush green foliage and trees. A small, dark, triangular sign is visible on a post in the background, partially obscured by the fallen trunk. The overall atmosphere is quiet and somewhat somber due to the fallen tree.

**If a link/LC/6pack fails
in a fabric...**


Optimizing for failure




Optimizing for failure



Facebook development cycle




Site Reliability Engineering



Development & production

- Keep development & production closely tied
- Developers have direct access to production environment
- On call
- Push out new code once a week

Summary



Build simple
reusable
components




Iterate fast
and make
them bullet
proof



Reuse to
build bigger
systems

For more information

- ➔ Introducing wedge - <https://goo.gl/YvNlwF>
- ➔ Operationalizing FBOSS - <https://goo.gl/VWLg3V>
- ➔ 6-Pack hardware - <https://goo.gl/umubpk>
- ➔ FBOSS on github - <https://github.com/facebook/fboss>



OPEN

Compute Project

