Evolution of Networks

Patrice Brissette

Distinguished Engineer

Powered by

SONiC

Evolution of Networks

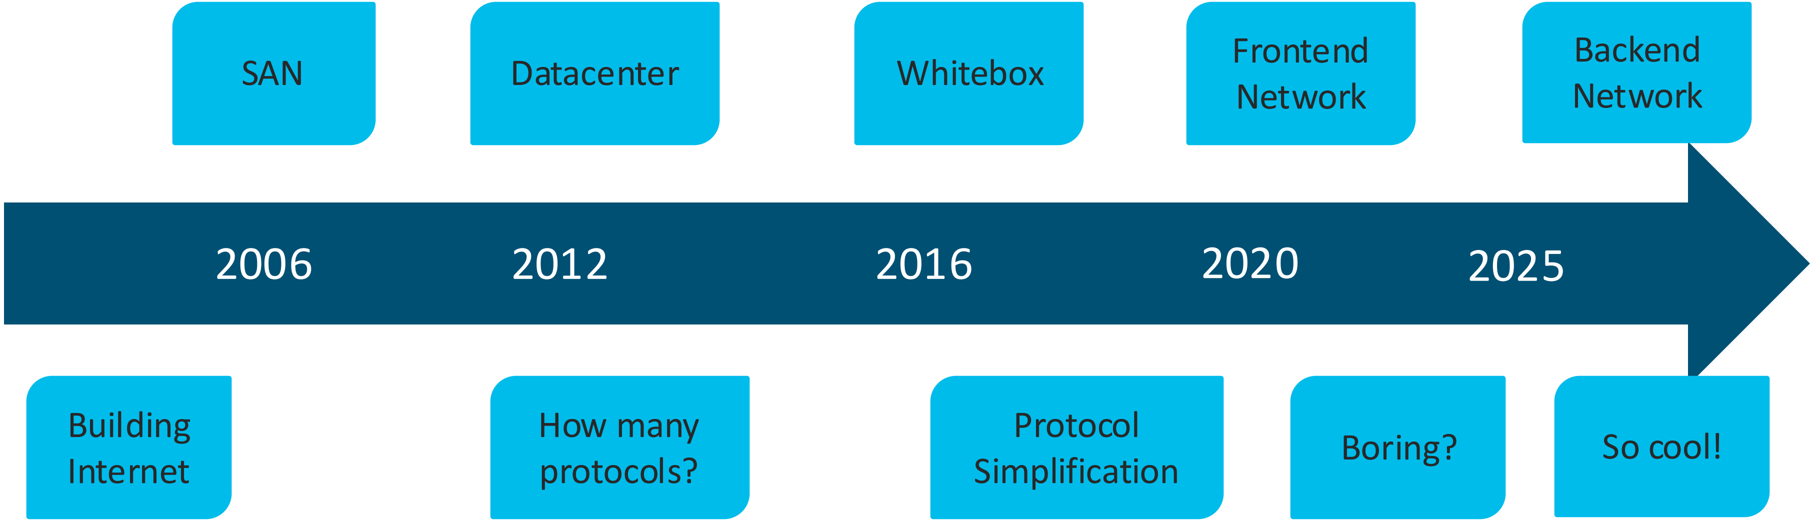

SAN

Datacenter

Whitebox

Frontend Network

Backend Network

2006

2012

2016

2020

2025

Building Internet

How many protocols?

Protocol Simplification

Boring?

So cool!

Network Representation

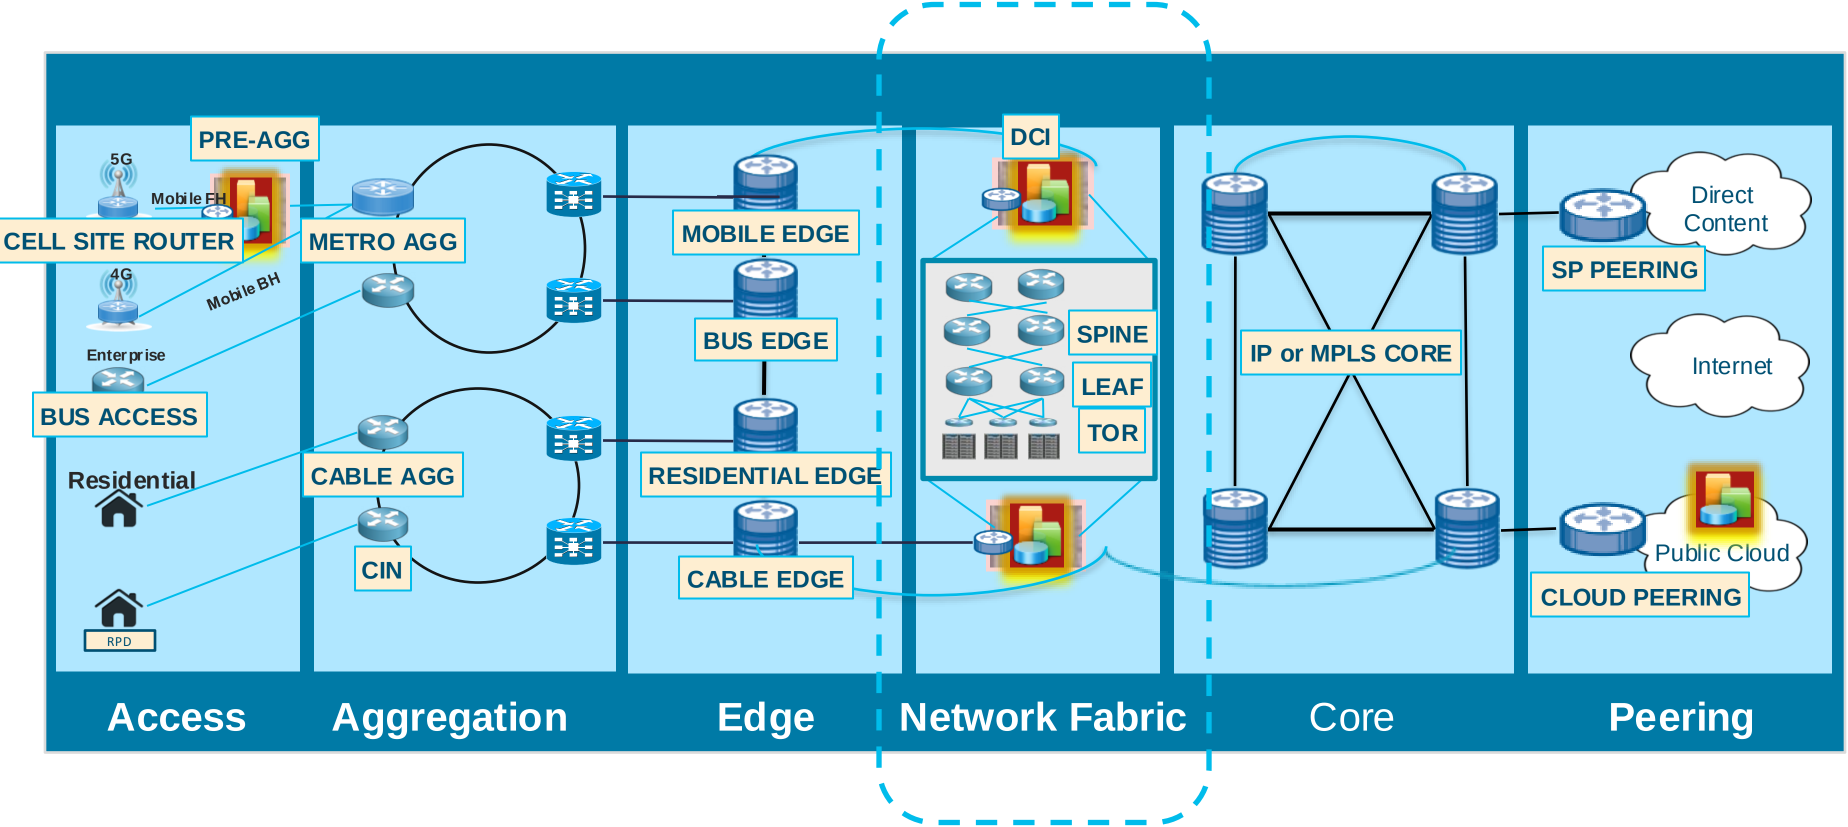

Front-End / Back-End Network

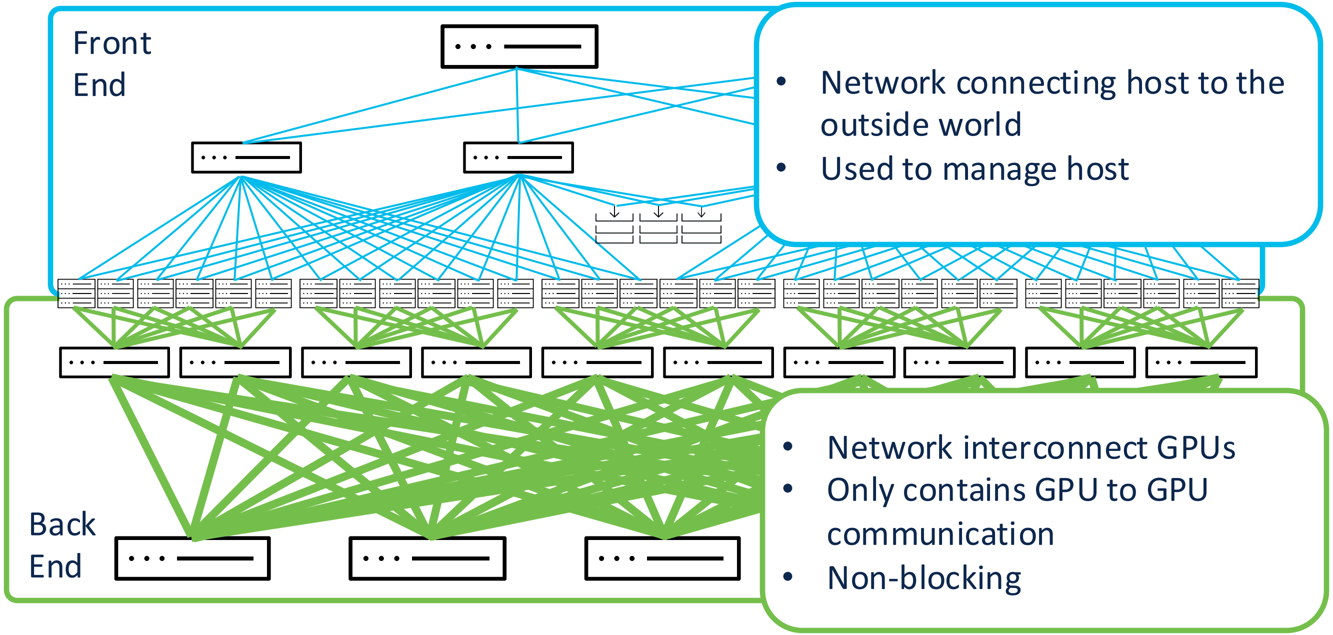

- Network connecting host to the outside world
- Used to manage host

- Network interconnect GPUs
- Only contains GPU to GPU communication
- Non-blocking

Scale-up: 4 to 8 GPU per server
Scale-out: servers are connected via proprietary protocol
Market is looking at options

Drastic Network Protocols Reduction

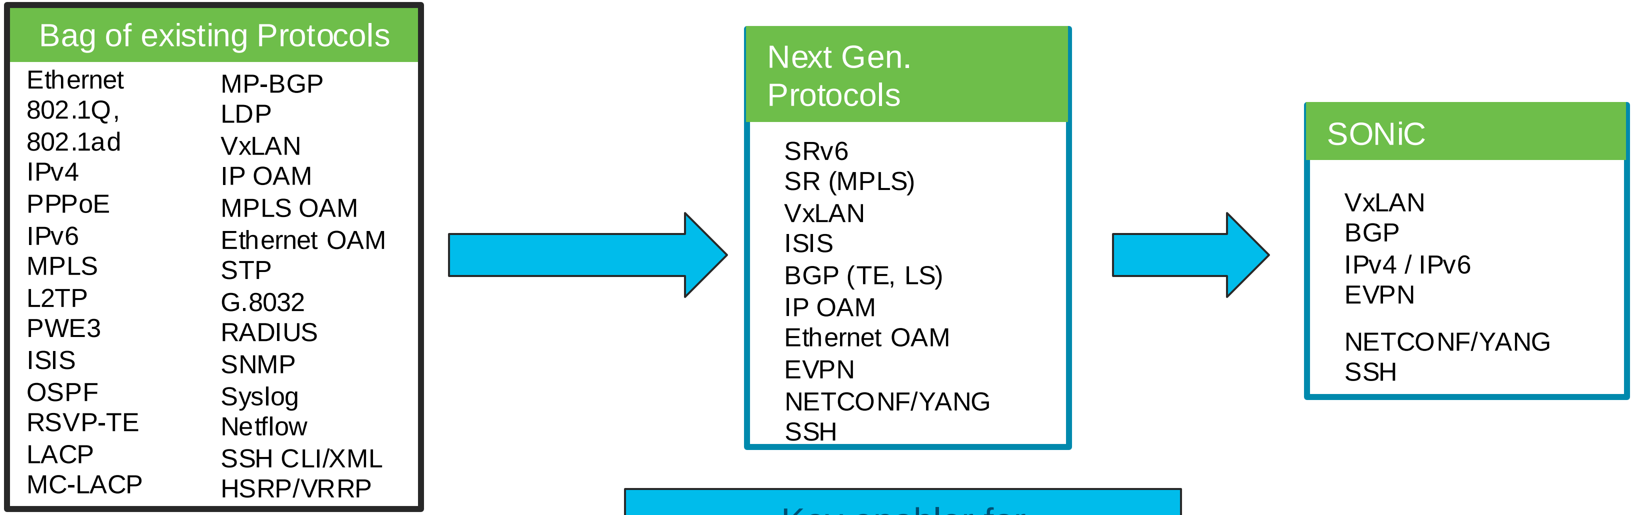

Bag of existing Protocols

Ethernet
802.1Q,
802.1ad
IPv4
PPPoE
IPv6
MPLS
L2TP
PWE3
ISIS
OSPF
RSVP-TE
LACP
MC-LACP

MP-BGP
LDP
VxLAN
IP OAM
MPLS OAM
Ethernet OAM
STP
G.8032
RADIUS
SNMP
Syslog
Netflow
SSH CLI/XML
HSRP/VRRP

Next Gen. Protocols

SRv6
SR (MPLS)
VxLAN
ISIS
BGP (TE, LS)
IP OAM
Ethernet OAM
EVPN
NETCONF/YANG
SSH

SONiC

VxLAN
BGP
IPv4 / IPv6
EVPN

NETCONF/YANG
SSH

Key enabler for

Reducing operations complexity

- Simpler automation
- Simpler to repair
- Simpler integration
- Foundation for service Orchestration

Requirements in 2020 Networks

Legacy Features	Support developed protocols since 1998 (> 40)
Convergence	50ms switch time upon **single** failure
Optimal Forwarding	Traffic engineering, proper bandwidth usage
Low Latency	5G requirements
ECMP	Entropy based on 5-tuple hash
Programmability	gRPC, gNMI, Yang, controller
Virtualization	Appliance running on server; disaggregation

Requirements in 2025 Networks

Non-Blocking	Throughput is critical for training purpose
Lossless	Congestion management

Bandwidth	Current networks to be enhanced (70% → 100%)
Low Latency	Inference network should be optimized for performance
Enhanced-ECMP	Elephant flow with low-entropy
Visibility	High-frequency Telemetry for network congestion control

SONiC

The SONiC's Journey

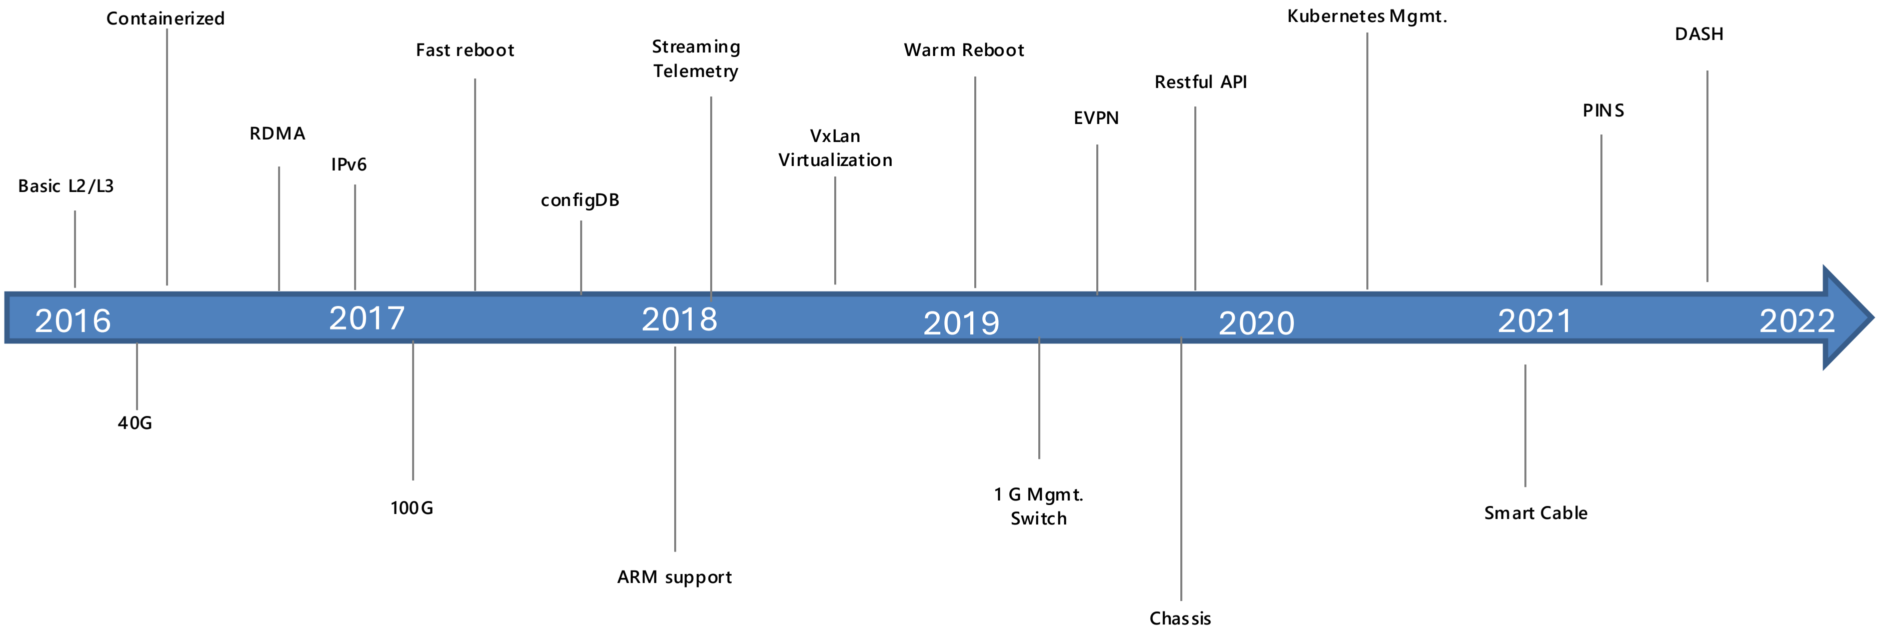

2016

Basic L2/L3

40G

Containerized

RDMA

2017

IPv6

100G

Fast reboot

configDB

ARM support

2018

Streaming Telemetry

VxLan Virtualization

2019

Warm Reboot

1 G Mgmt. Switch

EVPN

Chassis

Restful API

2020

Kubernetes Mgmt.

2021

Smart Cable

PINS

DASH

2022

SONiC: Software for Open Networking in the Cloud

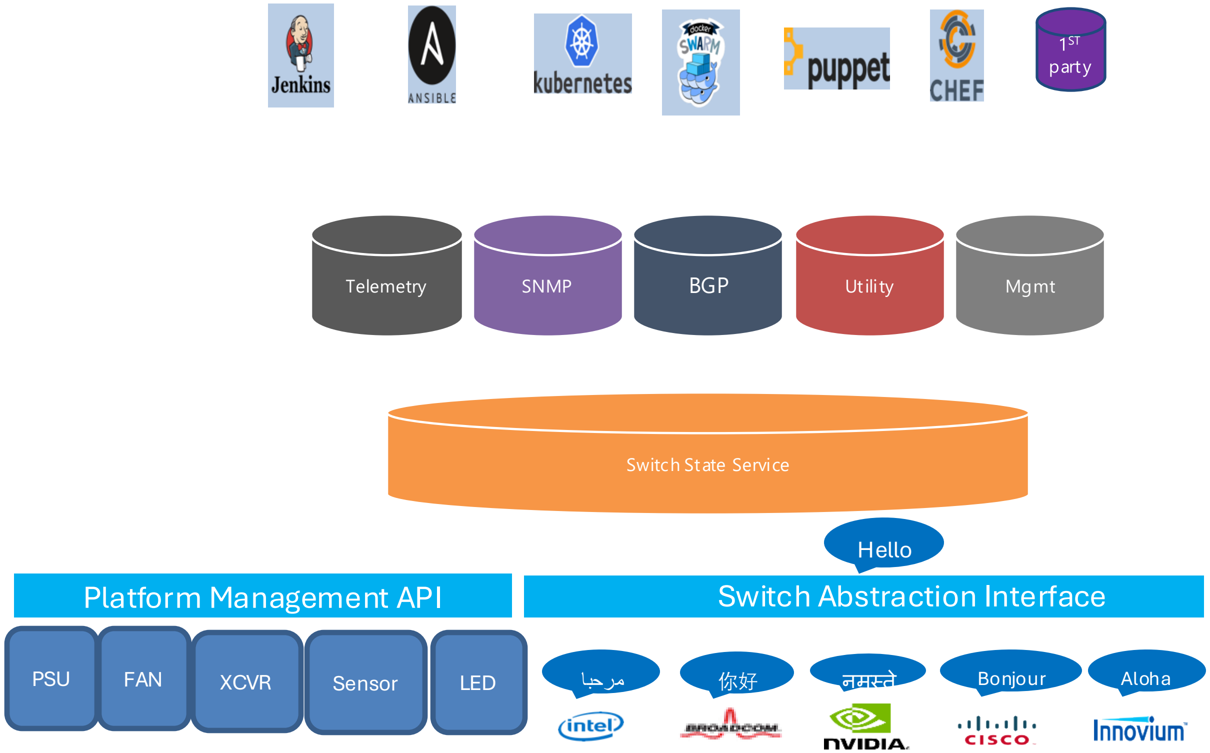

Rich management capabilities

Container for serviceability and portability

Application State Decoupling

Platform/ASIC Independence

Goals of SONiC

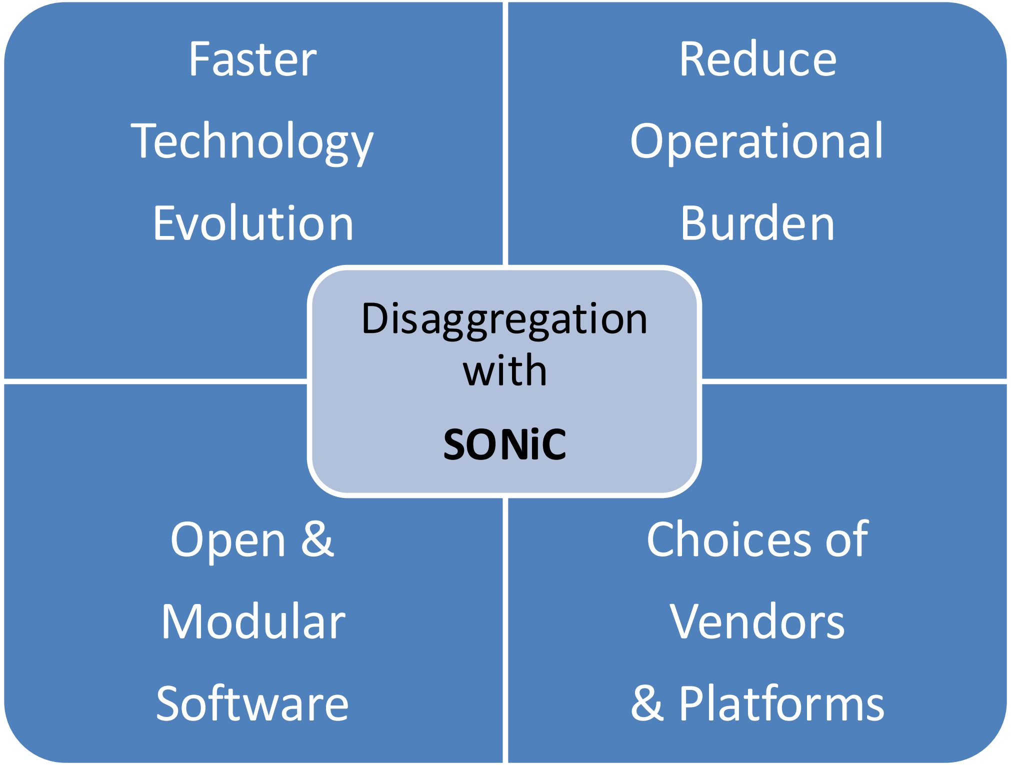

Faster Technology Evolution

Reduce Operational Burden

Disaggregation with **SONiC**

Open & Modular Software

Choices of Vendors & Platforms

SONiC/SAI Community and Partners

SONiC

Merchant Silicon

MEDIATEK
Nephos
CISCO
BROADCOM
MARVELL
Innovium
NVIDIA
intel
centec
AMD

Switch Platform

SUPERMICRO
H3C
RAGILE
ARISTA
CISCO
NOKIA
DELL
恒为科技 EmbedWay
Edge-corE NETWORKS
JUNIPER NETWORKS
Celestica
ALPHA Networks
Meta
NVIDIA
RAGILE
Quanta Computer
PEGATRON 和碩聯合科技
ufiSpace
DELTA
MiTAC
Inventec
inspur
CREDO we connect
WNC Wistron NeWeb Corp.
STORDIS

Adoption

Google
美团
Alibaba Group 阿里巴巴集团
Tencent 腾讯
DiDi
LinkedIn
JD Cloud
Microsoft
criteo
ebay
Bai du 百度
COMCAST
orange™
PANTHEON .tech

System/Service

DELL
apstra
wipro
BROADCOM
XCLOUD NETWORKS
Asterfusion
CANONICAL
KEYSIGHT TECHNOLOGIES
NVIDIA
香港科技大學 THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY
docker
insidepacket
AVIZ
augtera networks
HiWire ACTIVE ELECTRICAL CABLES
天翼云

SONiC - Year End Review

380K+

social media reach

27

SONiC members

4K

active contributors

6

2024 in-person events

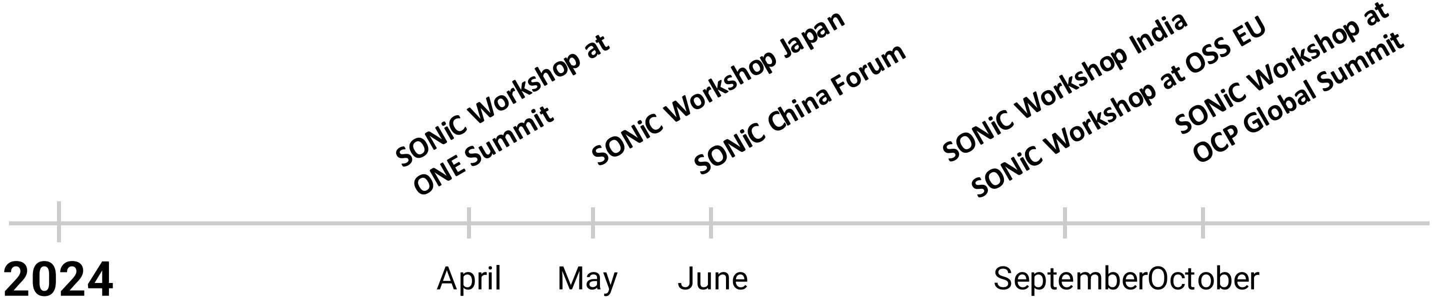

2024

April
SONiC Workshop at ONE Summit

May
SONiC Workshop Japan

June
SONiC China Forum

September
SONiC Workshop India
SONiC Workshop at OSS EU

October
SONiC Workshop at OCP Global Summit

Workgroups Highlights

AI Workgroup

- Fabric Workstream
  - Topology (led by Nvidia, Dell, BRCM, Marvell)
  - Routing (led by Nvidia, Dell, BRCM)
- Traffic Management Workstream
  - Congestion (led by Nvidia)
  - PFC (led by BRCM)
- Telemetry Workstream

(All workstreams finished first round investigation)

Smart Switch Workgroup

- Platform Architecture HLD for SmartSwitch completed
- DPU database design completed
- Midplane/PCIE architecture completed
- IP address assignment design completion
- gNMI design from NPU to DPU
- DPU High-availability design
- NPU-DPU Health probing design using BFD

OTN Workgroup

- OTN services (OTSS, Syncd-OT, OTAI) HLDs have been reviewed with PR to SONiC.
- OTN virtual platform and KVM device have been supported in SONiC-OTN branch.
- OTN enhancement on sonic-mgmt-common HLD has been reviewed with PR to SONiC.

Routing Workgroup

- Phoenix Wing initiated to integrate SRv6 features into SONiC platforms, enhancing compatibility and reducing total cost of ownership.
- Enhanced sonic-mgmt vsonic test infra from single box testing to multi-node topology testing.
- The phoenix wing platform test plan PR and the test infra enhancement code PR are raised.
- Spytest test plan is working in progress between Alibaba team and Cisco team.

Focus and Looking Forward

Security

- Vulnerability Report and Collaboration
- AI Empowered Security (Data Analyze, Threat Detection, Auto Repsonse etc)
- Secure Boot Enhancement
- Security Auditing & Insight
- Supply Chain Security

Ethernet for AI

- Optimized Topology
- Traffic management (Routing, Load Balance, Packet Spray etc)
- Congestion control
- Advanced telemetry
- Optics
- Liquid Cooling

Smart Switch

- Floating NIC
- Private Link
- Flow simulation
- Security Offloading
- Computation Offloading

Enterprise & Edge

- Power Over Ethernet (PoE)
- Multiple Spanning Tree Protocol (MSTP)

High-Level Architecture

High-Level Architecture

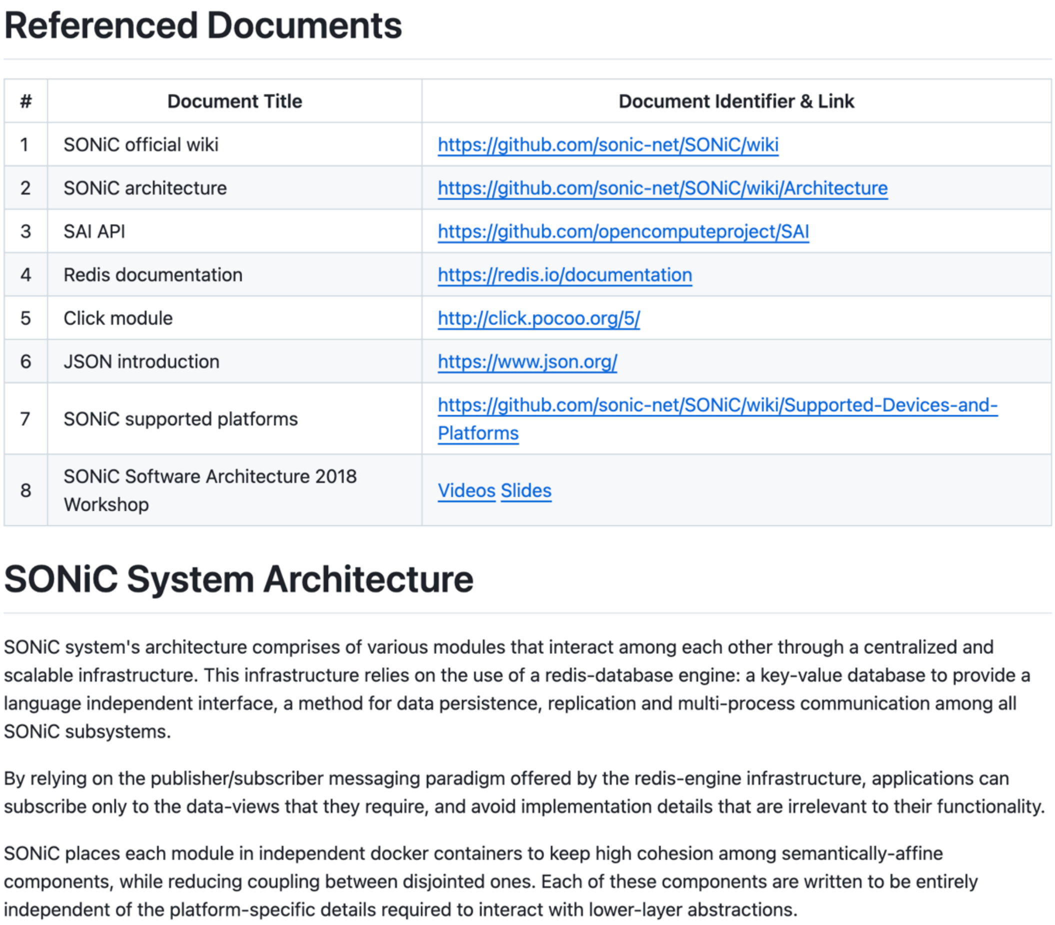

Referenced Documents

#	Document Title	Document Identifier & Link
1	SONiC official wiki	https://github.com/sonic-net/SONiC/wiki
2	SONiC architecture	https://github.com/sonic-net/SONiC/wiki/Architecture
3	SAI API	https://github.com/opencomputeproject/SAI
4	Redis documentation	https://redis.io/documentation
5	Click module	http://click.pocoo.org/5/
6	JSON introduction	https://www.json.org/
7	SONiC supported platforms	https://github.com/sonic-net/SONiC/wiki/Supported-Devices-and-Platforms
8	SONiC Software Architecture 2018 Workshop	Videos Slides

SONiC System Architecture

SONiC system's architecture comprises of various modules that interact among each other through a centralized and scalable infrastructure. This infrastructure relies on the use of a redis-database engine: a key-value database to provide a language independent interface, a method for data persistence, replication and multi-process communication among all SONiC subsystems.

By relying on the publisher/subscriber messaging paradigm offered by the redis-engine infrastructure, applications can subscribe only to the data-views that they require, and avoid implementation details that are irrelevant to their functionality.

SONiC places each module in independent docker containers to keep high cohesion among semantically-affine components, while reducing coupling between disjointed ones. Each of these components are written to be entirely independent of the platform-specific details required to interact with lower-layer abstractions.

Documents

High-Level Architecture

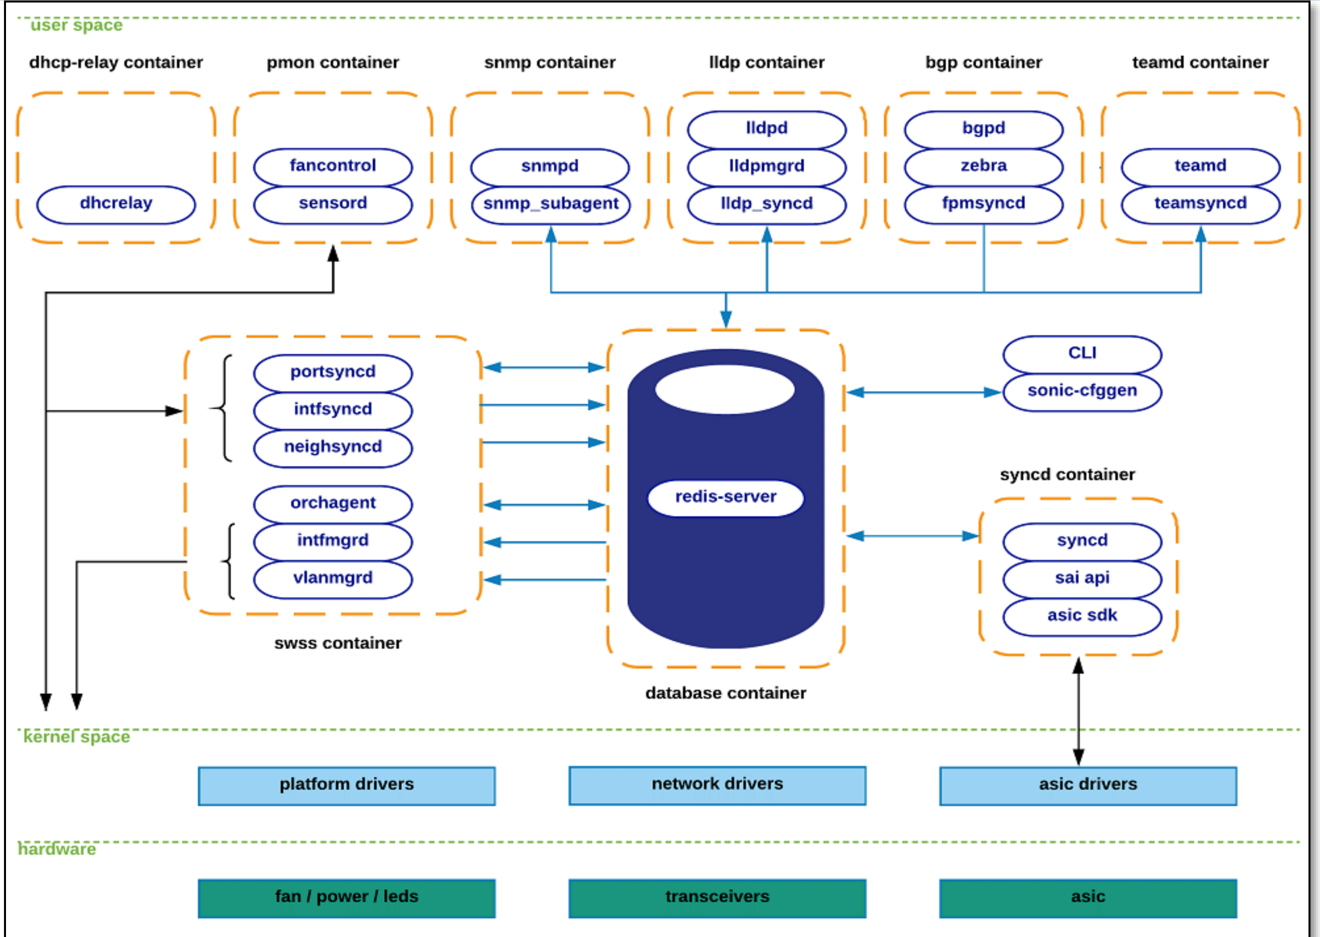

SONiC
Overview

High-Level Architecture

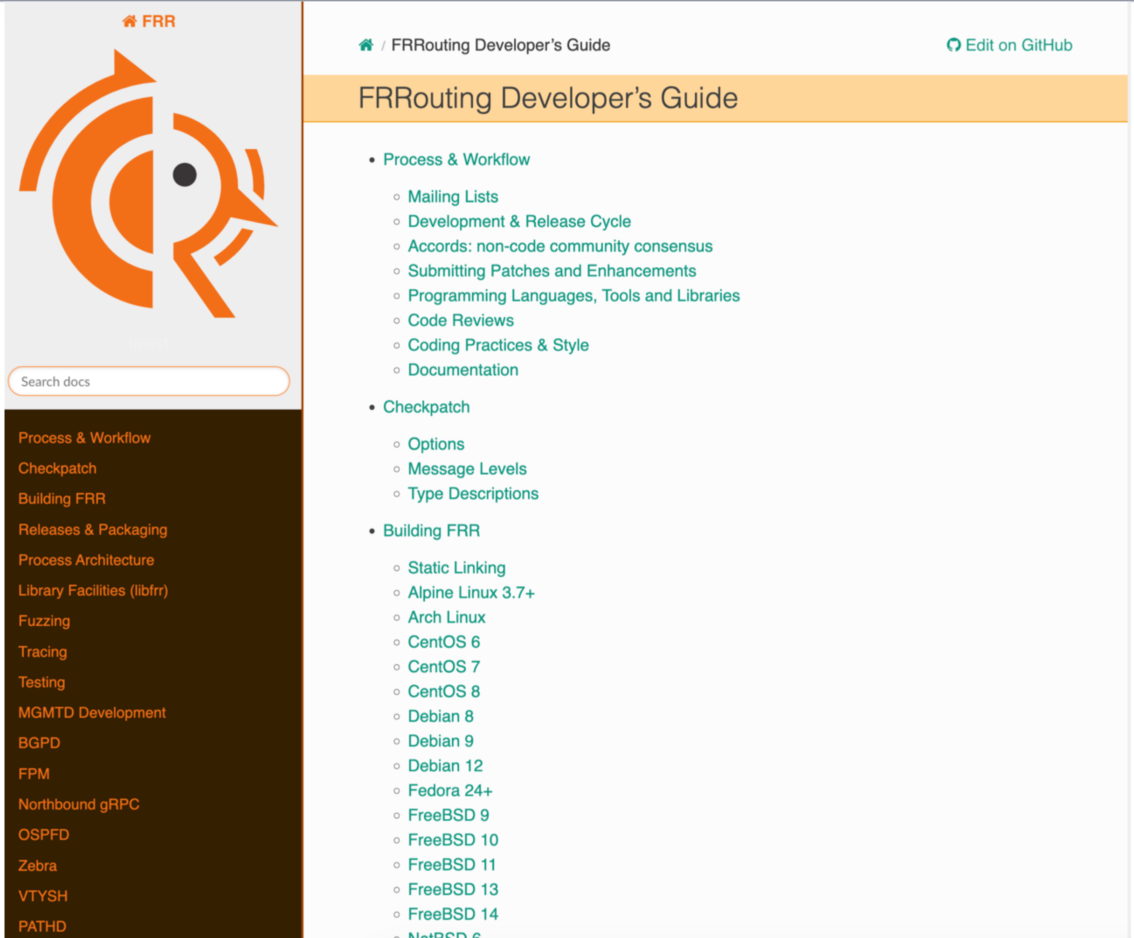

FRR

latest

Search docs

- Process & Workflow
- Checkpatch
- Building FRR
- Releases & Packaging
- Process Architecture
- Library Facilities (libfrr)
- Fuzzing
- Tracing
- Testing
- MGMTD Development
- BGPD
- FPM
- Northbound gRPC
- OSPFD
- Zebra
- VTYSH
- PATHD

/ FRRouting Developer's Guide

Edit on GitHub

FRRouting Developer's Guide

- Process & Workflow
  - Mailing Lists
  - Development & Release Cycle
  - Accords: non-code community consensus
  - Submitting Patches and Enhancements
  - Programming Languages, Tools and Libraries
  - Code Reviews
  - Coding Practices & Style
  - Documentation
- Checkpatch
  - Options
  - Message Levels
  - Type Descriptions
- Building FRR
  - Static Linking
  - Alpine Linux 3.7+
  - Arch Linux
  - CentOS 6
  - CentOS 7
  - CentOS 8
  - Debian 8
  - Debian 9
  - Debian 12
  - Fedora 24+
  - FreeBSD 9
  - FreeBSD 10
  - FreeBSD 11
  - FreeBSD 13
  - FreeBSD 14
  - NetBSD 6

FRRouting

Useful Links

https://github.com/sonic-net/

https://github.com/sonic-net/SONiC/wiki

https://github.com/sonic-net/SONiC/wiki/Architecture

https://github.com/opencomputeproject/SAI

https://frrouting.org/

https://docs.frrouting.org/projects/dev-guide/en/latest/index.html#

Overview

Started in May 2023.

Charter:
Gather contributors from both FRR community and SONiC community to improve white box support for a router based on deployments / use cases.

Chair:	Eddie Ruan (Alibaba)
Co-chairs:	Patrice Brissette (Cisco) Syed Hasan Raza Naqvi (Broadcom) Rita Hui (Microsoft)
Webpage:	https://lists.sonicfoundation.dev/g/sonic-wg-routing
Mailing List:	sonic-wg-routing@lists.sonicfoundation.dev
Subscribe:	sonic-wg-routing+subscribe@lists.sonicfoundation.dev
Members:	300+

Weekly Meeting:
Every Wed, 11am – 12pm EST
Every Thu, 8.30pm – 9.30pm EST

Workflow

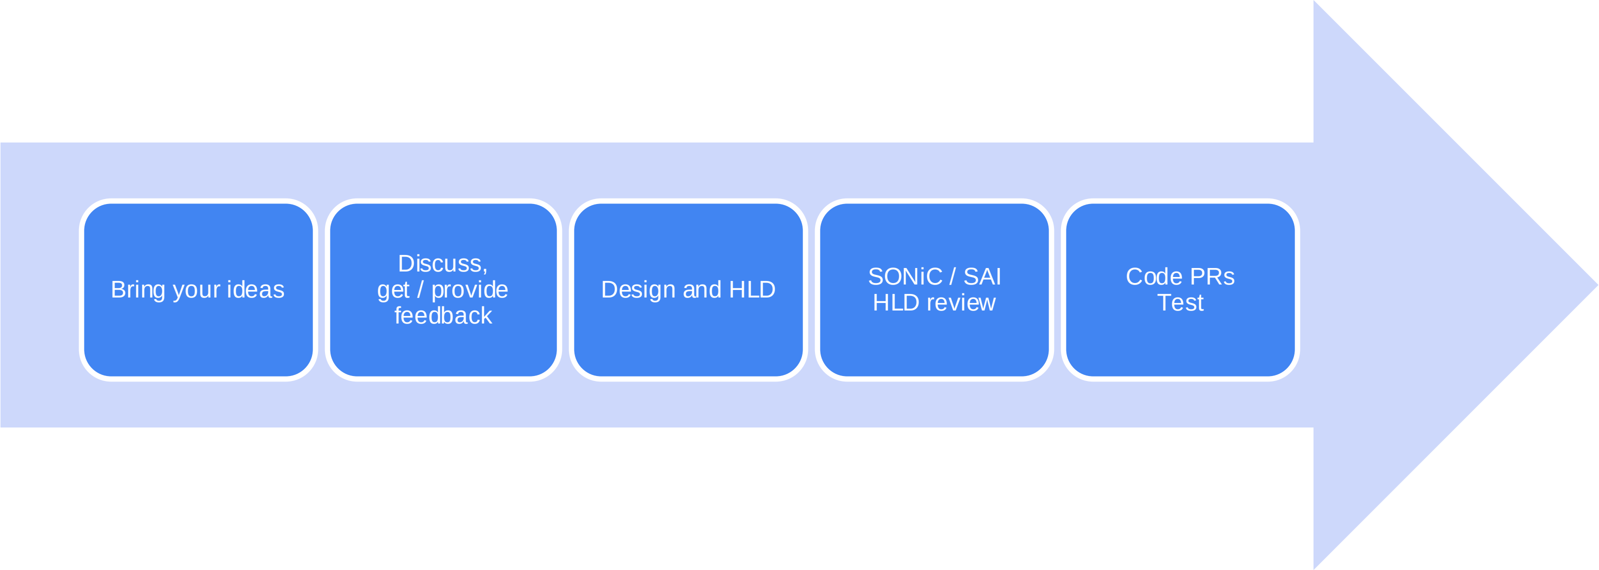
Bring your ideas
Discuss, get / provide feedback
Design and HLD
SONiC / SAI HLD review
Code PRs Test

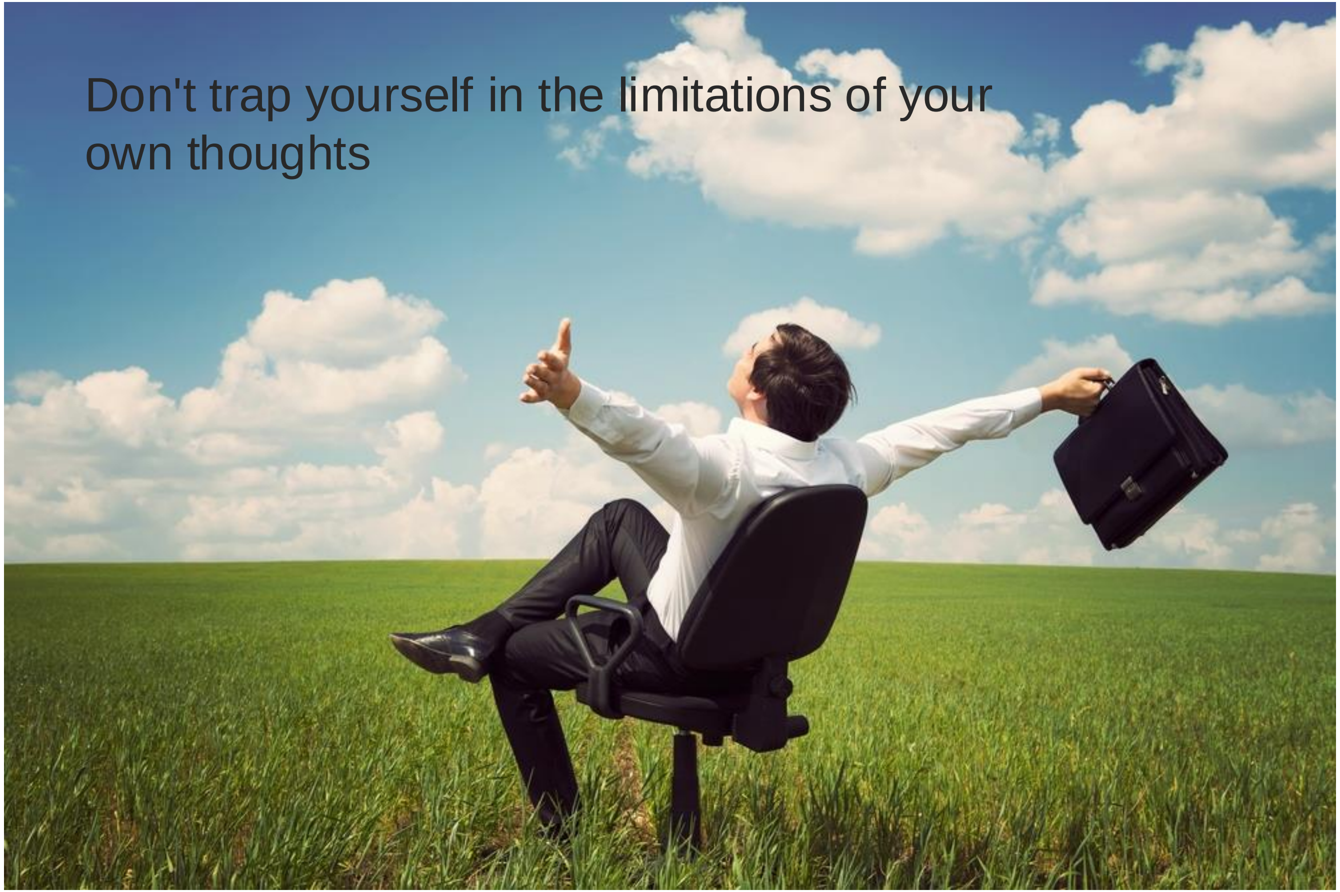

Don't trap yourself in the limitations of your own thoughts

How to reach me?

pbrisset@cisco.com

Alamy