



OPEN
Compute Project®

SONiC Update

OCP Workshop August, 2018

Xin Liu, Microsoft

SONiC Recap – Platform Agnostic NOS



SAI

Hardware



Network Applications

Hello




Switch Abstraction Interface

частный

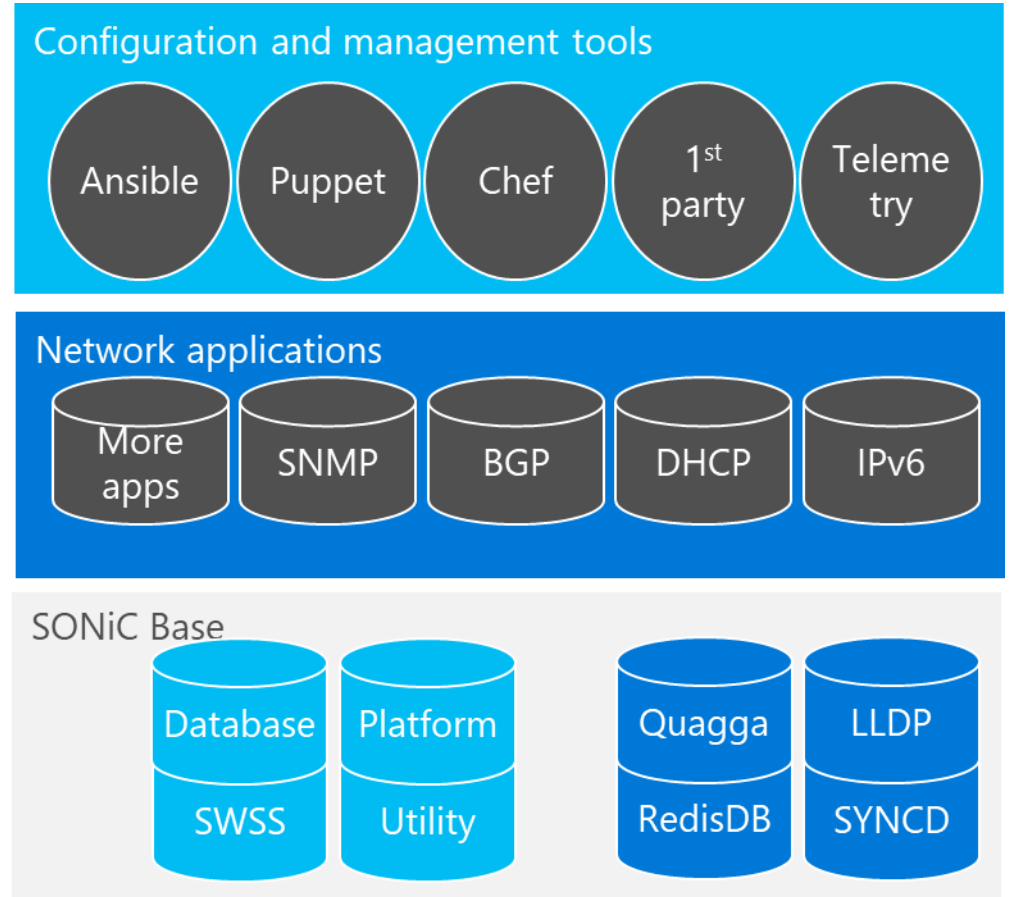

你好

नमस्ते

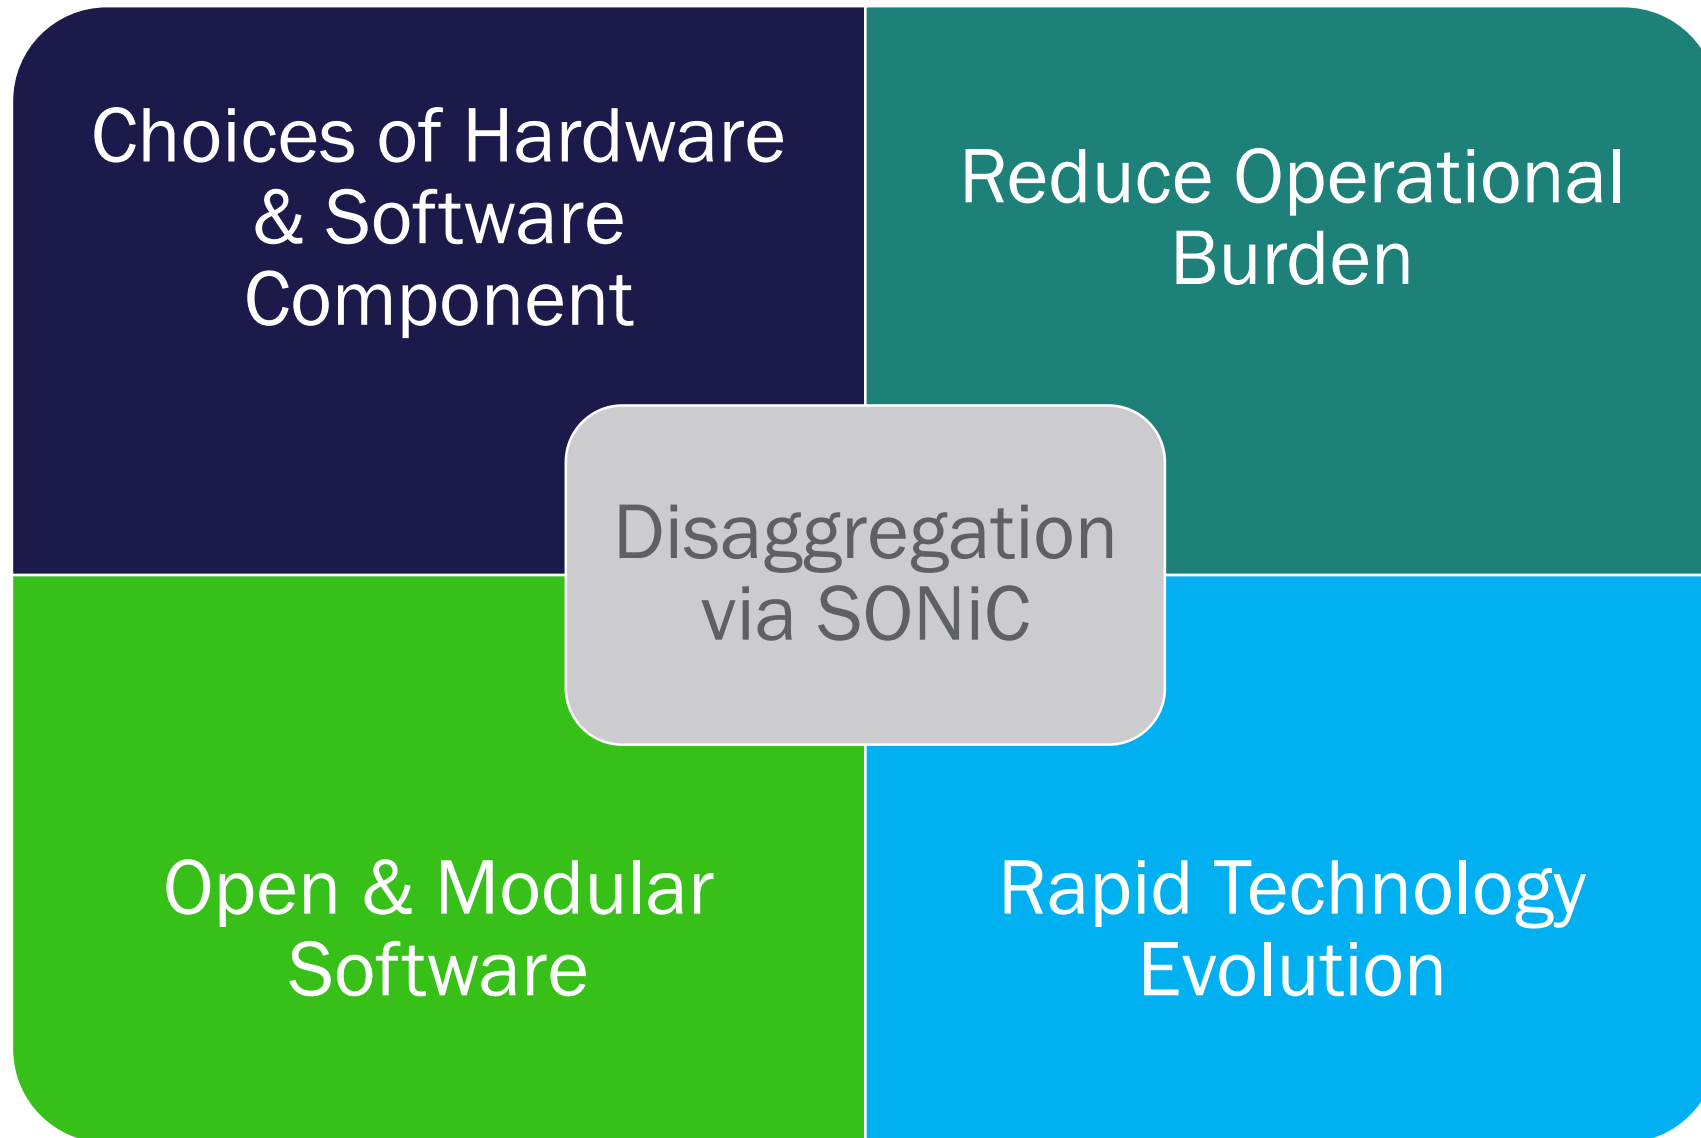
Bonjour




SONiC Recap – Containerized NOS



SONiC Recap – Open Source NOS Evolving Fast



SONiC Releases



What Has Been Worked On?


Release	Release Date	SAI version	Features Included
SONiC.201803	3/30/2018	1.2	Critical Resource Monitoring
			MAC Aging
			IPv6 ACL
			BGP/Neighbor-down fib-accelerate
			PFC WD

Release	Release Date	SAI version	Features Included
SONiC.201807	7/30/2018	1.3	gRPC
			Dtel support
			SONiC Architecture and User Manual (Documentation), will be updated soon
			Sensor transceiver monitoring
			LLDP extended MIB

Release	Release Date	SAI version	Features Included
SONiC.201810	10/15/2018	1.3	Debian Kernel Upgrade to 4.9
			Warm Reboot
			CLI Framework
			BGP Graceful Restart
			Incremental Config (IP, LAG, Port shut/unshut)
			Asymmetric PFC
			PFC Watermark

What is PFC?

- **PFC** (priority-based flow control) -> **lossless traffic** -> RDMA




Upstream egress queue stuck

Upstream neighbor keeps flooding PFC if its ingress queue reaches pause threshold

Problem: NIC PFC pause frame storm

- Malfunction NIC:
 - Keep sending PFC pause frames to ToR
- ToR/T1/T2 flood PFC pause to all neighbors



The whole network blocks by a single mad NIC

Related Developments

- PFC WD - Prevent single mad server from blocking the whole network
 - Detect the pause frame storm
 - Recover the storm after detection.
- Asymmetric PFC – Interface honors pause frame from all priorities, but only generates pause frame on lossless priority
 - To handle problem in the scenario that NIC queues all traffic to one queue
- Headroom Watermark – Troubleshooting enhancement

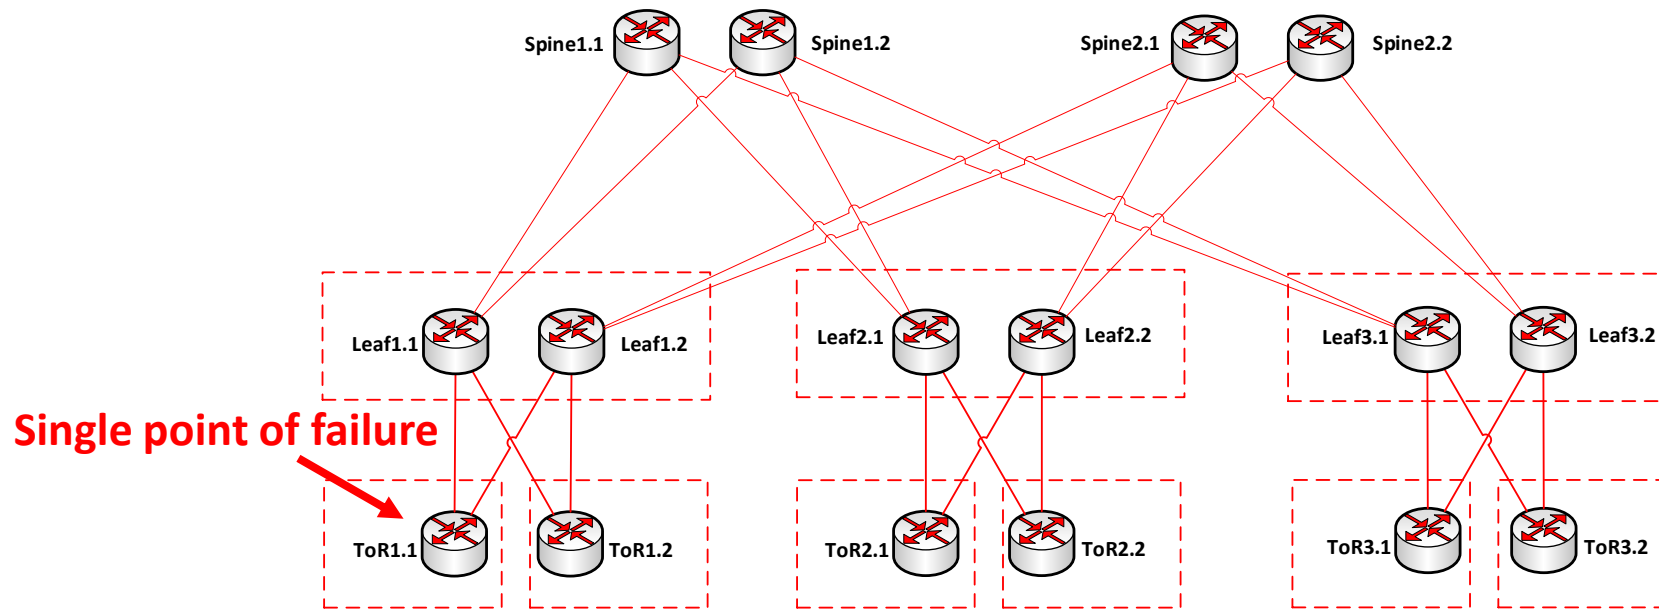
What Has Been Worked On?

Release	Release Date	SAI version	Features Included
SONiC.201803	3/30/2018	1.2	Critical Resource Monitoring
			MAC Aging
			IPv6 ACL
			BGP/Neighbor-down fib-accelerate
			PFC WD

Release	Release Date	SAI version	Features Included
SONiC.201807	7/30/2018	1.3	gRPC
			Dtel support
			SONiC Architecture and User Manual (Documentation), will be updated soon
			Sensor transceiver monitoring
			LLDP extended MIB

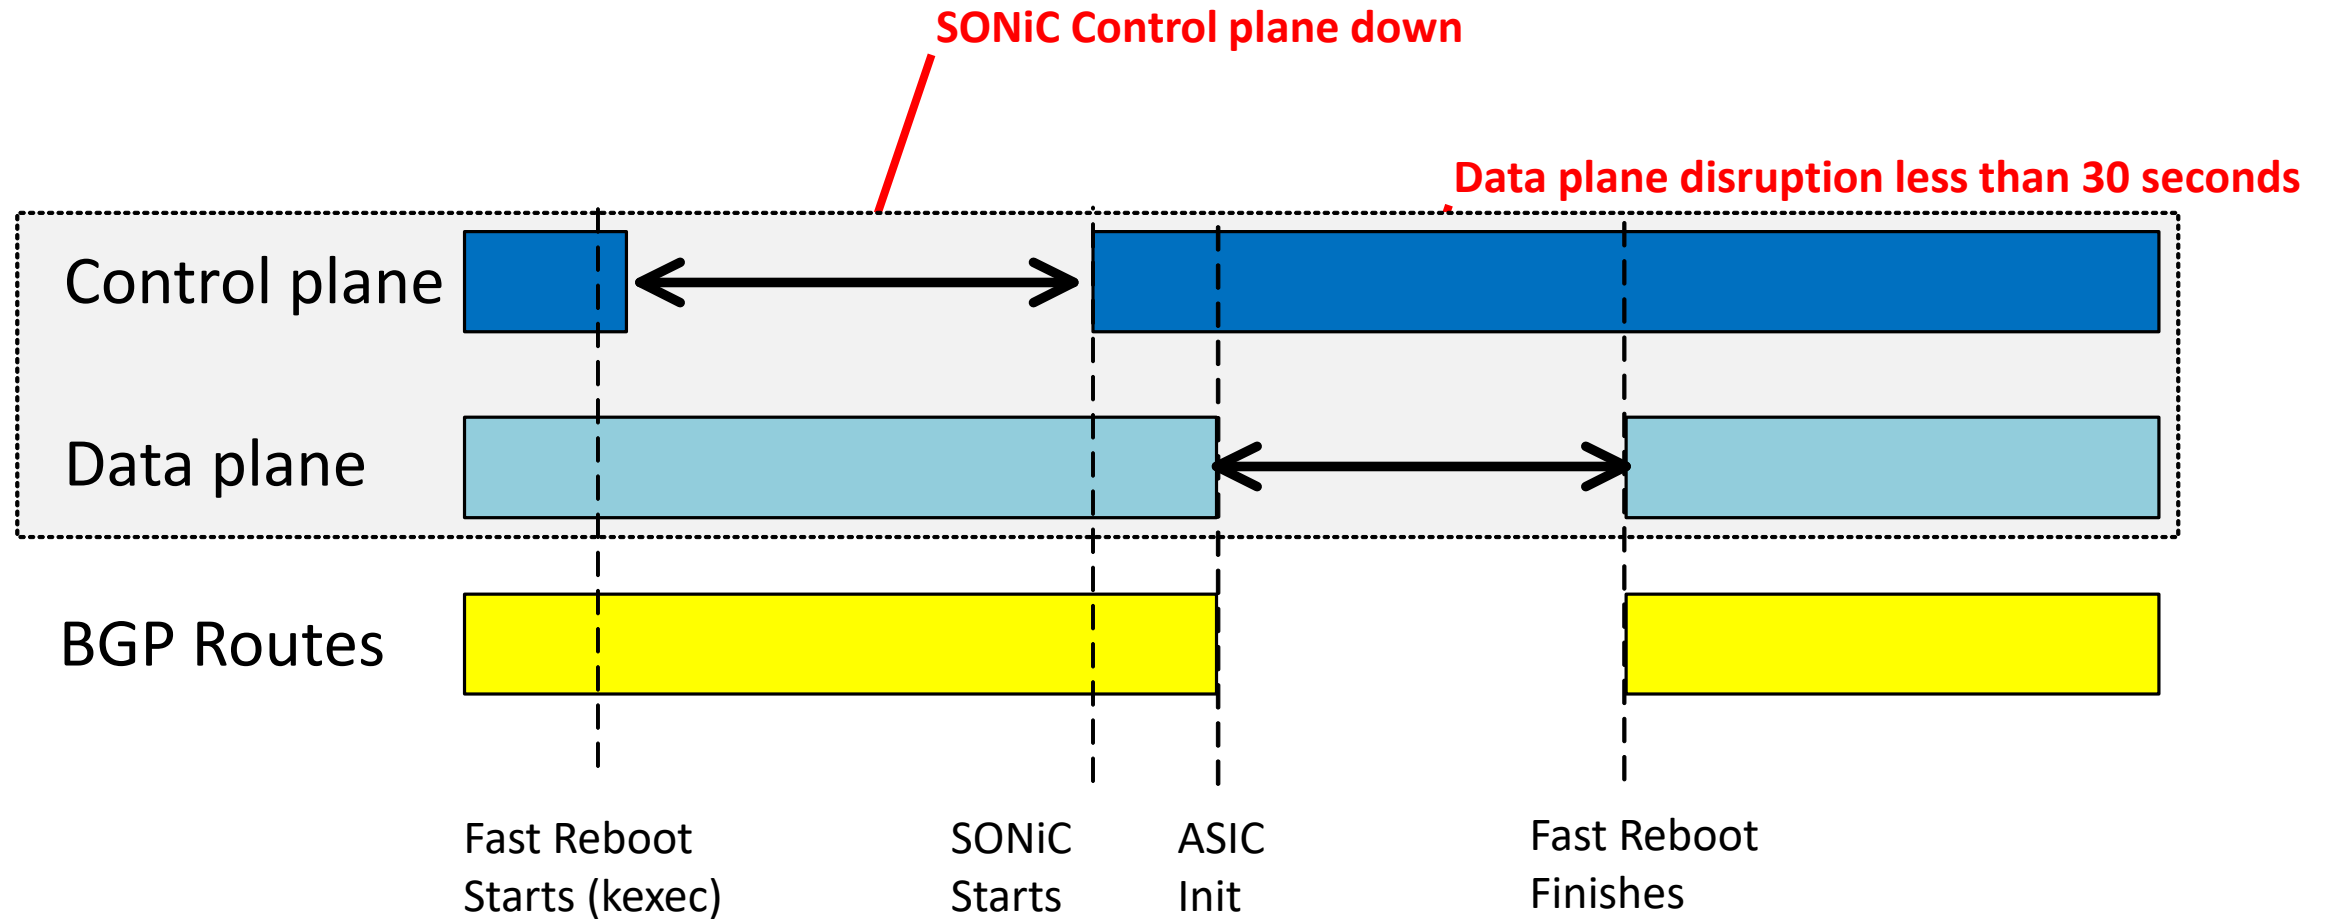
Release	Release Date	SAI version	Features Included
SONiC.201810	10/15/2018	1.3	Debian Kernel Upgrade to 4.9
			Warm Reboot
			CLI Framework
			BGP Graceful Restart
			Incremental Config (IP, LAG, Port shut/unshut)
			Asymmetric PFC
			PFC Watermark

Velocity + Serviceability + Reliability

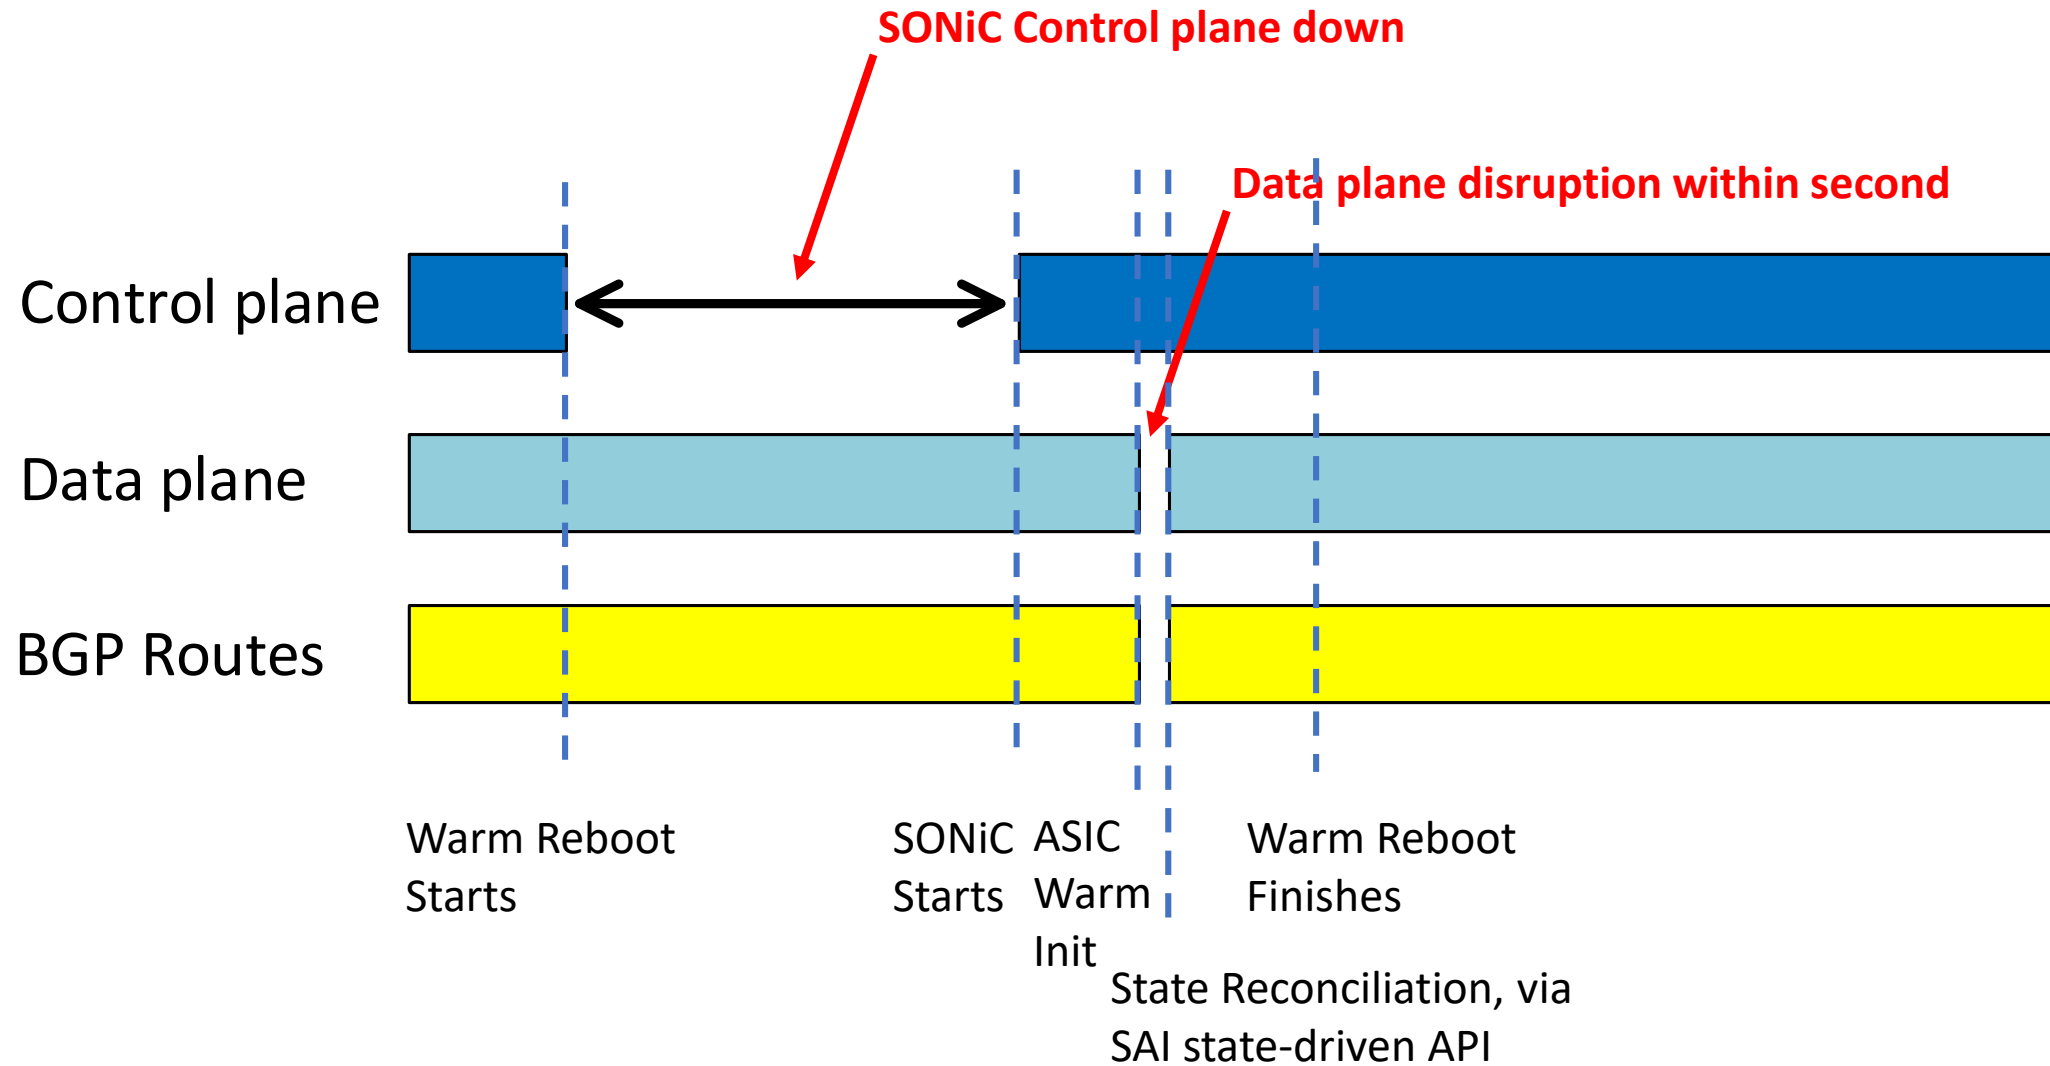


- Critical customer applications have “zero tolerance” to the network disruption, e.g., game streaming, AI
- Continuously upgrade software on the TOR is needed to fix bugs and add features

Fast-reboot Upgrade



Warm reboot Upgrade



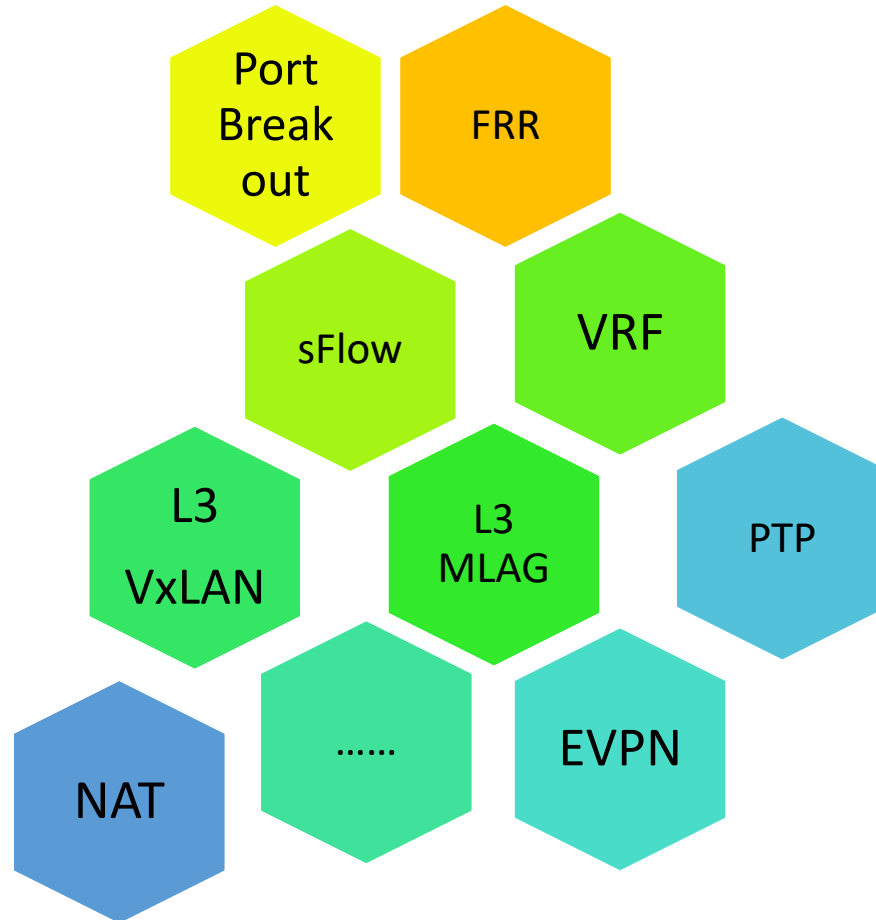
What Has Been Worked On? + New ASIC + New Platform

Release	Release Date	SAI version	Features Included
SONiC.201803	3/30/2018	1.2	Critical Resource Monitoring
			MAC Aging
			IPv6 ACL
			BGP/Neighbor-down fib-accelerate
			PFC WD

Release	Release Date	SAI version	Features Included
SONiC.201807	7/30/2018	1.3	gRPC
			Dtel support
			SONiC Architecture and User Manual (Documentation), will be updated soon
			Sensor transceiver monitoring
			LLDP extended MIB

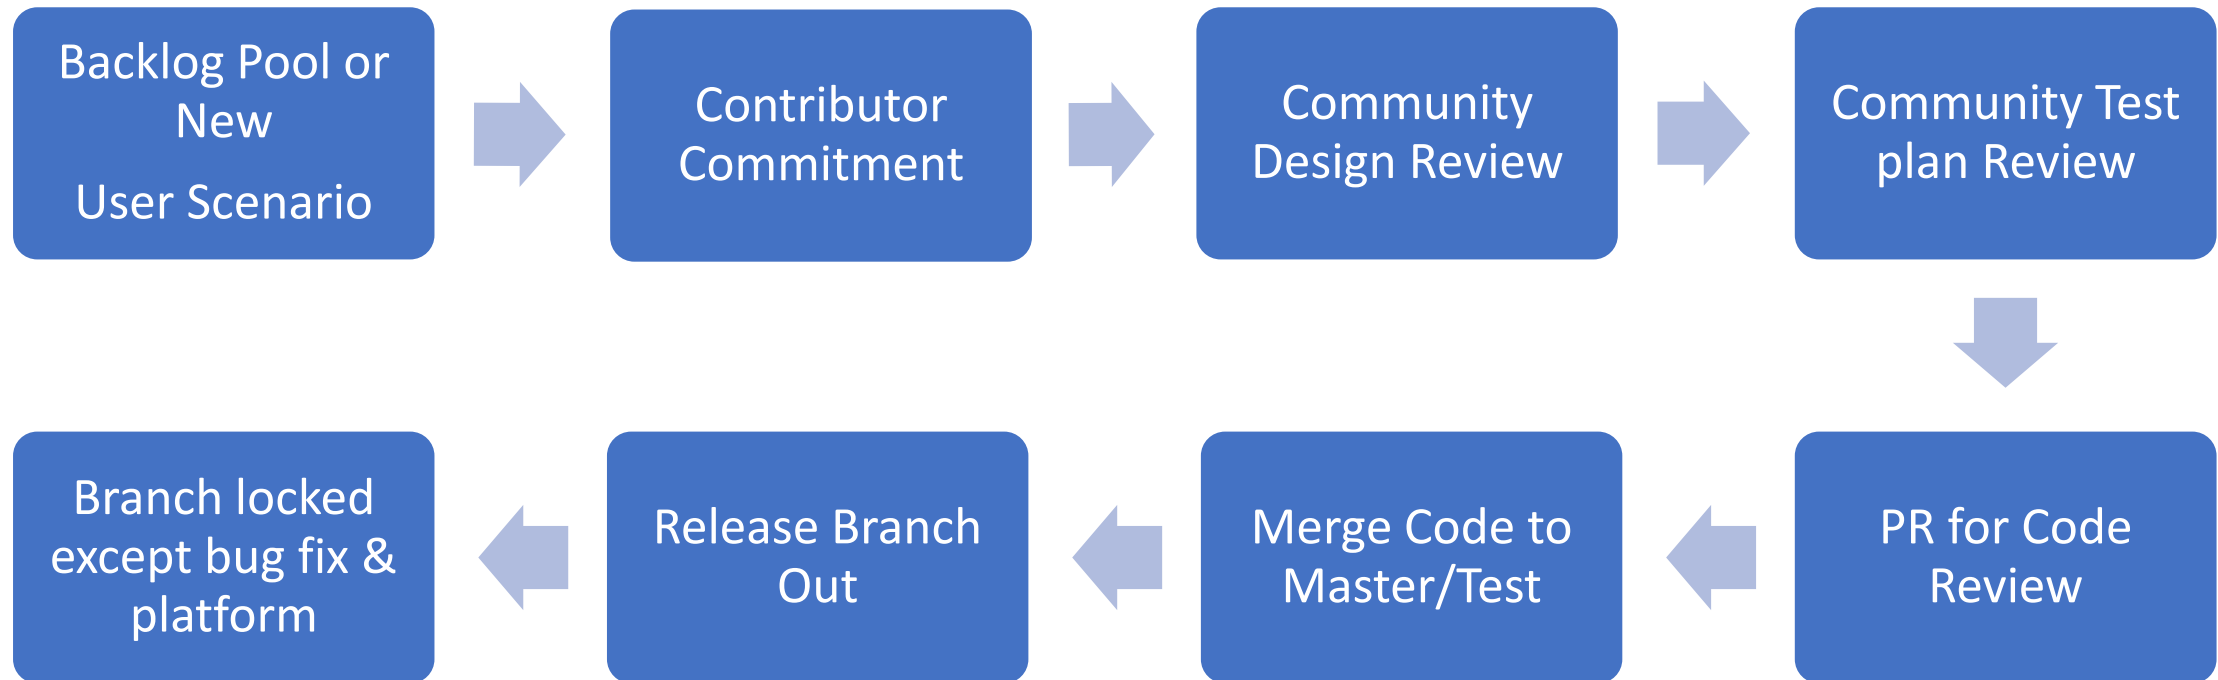
Release	Release Date	SAI version	Features Included
SONiC.201810	10/15/2018	1.3	Debian Kernel Upgrade to 4.9
			Warm Reboot
			CLI Framework
			BGP Graceful Restart
			Incremental Config (IP, LAG, Port shut/unshut)
			Asymmetric PFC
			PFC Watermark

Looking Forward



<https://github.com/Azure/SONiC/wiki/Sonic-Roadmap-Planning>

How Does a Release Work?



Community Activities

- Community Meeting
 - Calendar: <https://www.opencompute.org/projects/networking>
 - Time: Bi-weekly Tuesday 8am ~ 9am PST
 - Recording posted at Youtube (<https://www.opencompute.org/wiki/Networking/SONiC>)
- Mailing List
 - sonicproject@googlegroups.com
- Issue Tracking
 - <https://github.com/Azure/SONiC/issues>
- Slack Discussion Group
 - sonicswitch.slack.com

Afternoon Breakout Session

Time	Theme	Talks
1:30 ~ 2:30 pm	Build, Test	<ul style="list-style-type: none">- Platform Porting Guide [Joe from MSFT]- Building Barefoot image [Arkadiy from Barefoot]- Run SAI test on PTF [Zhichen from MLNX]
15 mins break		
2:45 ~ 3:45 pm	Tutorial [Rodny from LNKD]	<ul style="list-style-type: none">- SONiC Architecture- Case Study<ul style="list-style-type: none">• How to enable FRR on SONiC?• Upgrade a docker in SONiC image• How to integrate FRR into SONiC?- Example to enable people to add CLI
15 mins break		
4:00 ~ 4:30 pm	Telemetry [Jipan from Alibaba, Zhenggen from LNKD]	<ul style="list-style-type: none">- Telemetry arch and benefit- Demo
4:30 ~ 5:00 pm	Q & A	

Backup