Design Driven Network Assurance

Jeremy Schulman @ NANOG-88

Agenda

- Background and Motivation
- Demonstration
- Design Composition
- Validating a Network
- Q & A

My hope for your takeaway:

Consider different approach towards network automation and Infrastructure as Code

Spark ideas how you approach your next project

Background and Motivation

- Any network design
- Any vendors
- Greenfield projects
- Brownfield projects
- Agile

- Multi-site enterprise network
- Multiple network vendors
- Multiple networking domains
- Complex network edge requirements
- Quickly adapt to "in field" project changes
- Challenging timelines due to external factors

Network Quality Assurance

Ensure the network is operating as expected and report any anomalies **with as much context as possible maximizing situational awareness**.

For any given network we define a ***design*** from which we automatically derive ***checks.*** We use these checks to validate the design against the current state of the network.

The checks are automatically recomputed when we make changes to the design

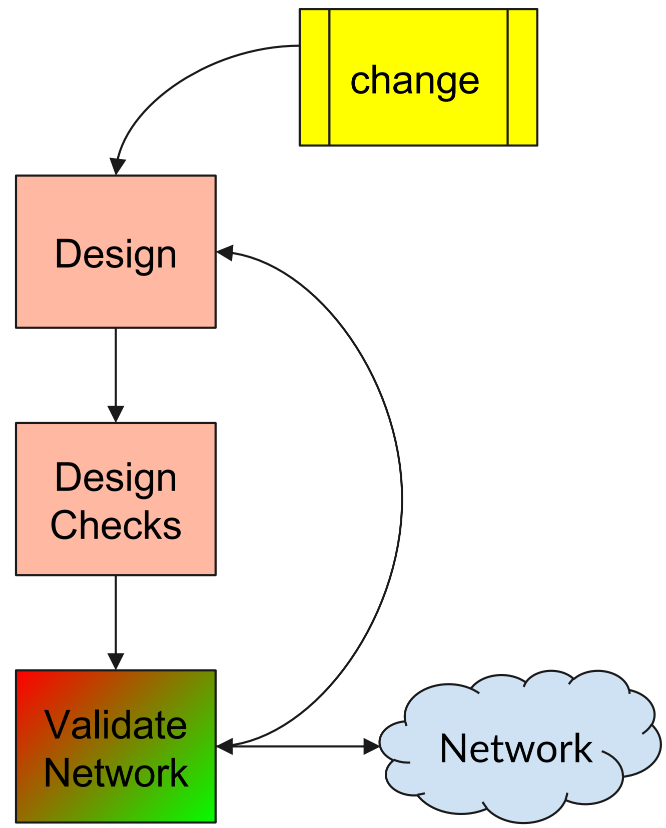

Quick Demo

Executing the design checks against the live network and reporting network assurance results

Device: sp01.swb.aaa, Total Results: 225

Test Cases	Status	Total	Pass	Fail	Info	Skip
ptp-system	FAILED!	1	0	1	0	0
ptp-ports	FAILED!	29	26	3	0	0
device	PASSED!	2	1	0	1	0
interfaces	FAILED!	77	73	4	0	0
transceivers	FAILED!	37	36	1	0	0
cabling	PASSED!	8	8	0	0	0
lags	PASSED!	3	3	0	0	0
ipaddrs	PASSED!	17	17	0	0	0
vlans	PASSED!	18	18	0	0	0
switchports	PASSED!	33	33	0	0	0

Network "Sources of Truth"

- The current operating state of the network is a source of truth
- A network design is a source of truth
    - A design may not encompass every aspect of the operating state
    - 80/20 rule and focus on most valued conditions
- A network assessment is the comparison of these two truths
    - Any difference represents an *anomaly*
    - Could require a change in the design
    - Could require a change in network config(s)
    - Could require a "fix" such as replacing a transceiver

What About Network Config Management?

- Network configurations are **NOT** a source of truth
  - They affect the operating state
  - By-product derived from design
- Real-world constraints
  - Manual changes by humans - generally "break-fix"
  - Automated changes by tools - different methods for generating and applying configurations

Design Composition

Network Infrastructure as Code

- What is a Design
- Design "Compiler"
- Design Reports
- Design Services

Design "Compiler"

- Builds the design and reports design errors
- Generates the checks for each device, executed later by a validation engine
- Generates reports
- Optionally generates device configurations

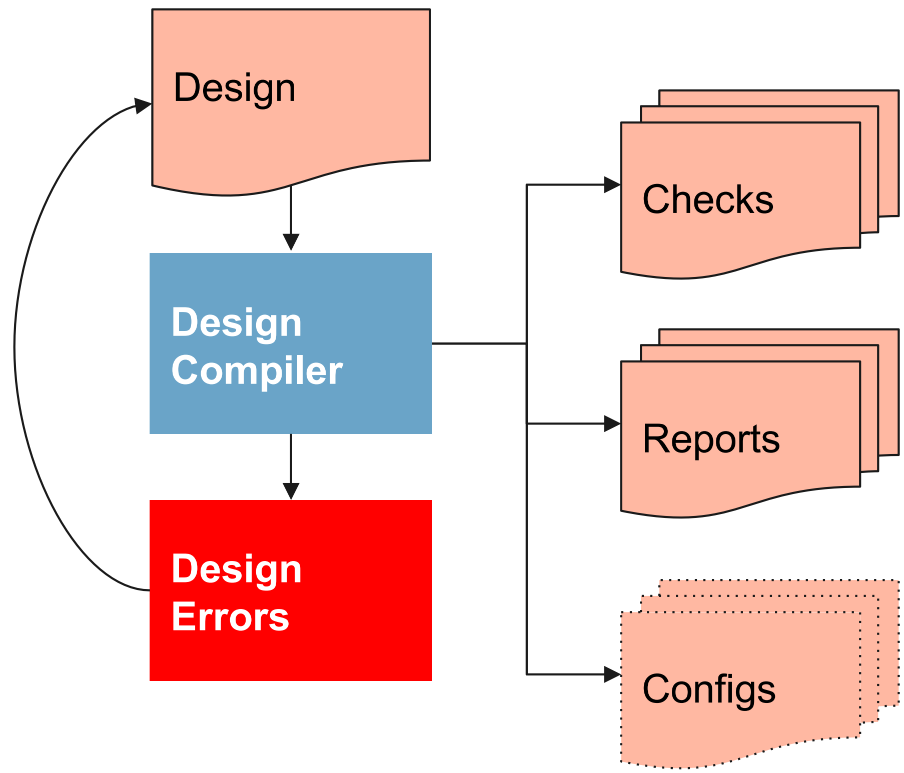

What is a Design?

The set of devices in a given networked together

Device	Profile	OS	Product Model	Primary IP	Managed Mode
lf01.swb.aaa	Leaf01Device	[redacted]	[redacted]	[redacted]	complete
lf02.swb.aaa	Leaf02Device	[redacted]	[redacted]	[redacted]	complete
sp01.swb.aaa	SpineDevice	[redacted]	[redacted]	[redacted]	complete
string01.swb.aaa	StringerDevice	[redacted]	[redacted]	[redacted]	complete
tr01.swb.aaa	TransitDevice	[redacted]	[redacted]	[redacted]	complete

The composition of *design services* that each device uses to validate the state of the network

Service Name	Kind	Checks	Devices
ptp	PTPDesignService	ptp-ports, ptp-system	lf01.swb.aaa, lf02.swb.aaa, sp01.swb.aaa, string01.swb.aaa
topology	TopologyDesignService	cabling, device, interfaces, ipaddrs, lags, transceivers	lf01.swb.aaa, lf02.swb.aaa, sp01.swb.aaa, string01.swb.aaa, tr01.swb.aaa
vlans	VlansDesignService	switchports, vlans	lf01.swb.aaa, lf02.swb.aaa, sp01.swb.aaa, string01.swb.aaa, tr01.swb.aaa

Network Quality Assurance Matrix

	Topology	VLANs	PTP	BGP-Peer	Multicast
dev1					
dev2					
dev3					
dev4					
dev5					
dev6					
dev7					

Network Check Results … Pass/Fail

	Topology	VLANs	PTP	BGP-Peer	Multicast
dev1					
dev2					
dev3					
dev4					
dev5					
dev6					
dev7					

"dev3" is Operating as Expected

	Topology	VLANs	PTP	BGP-Peer	Multicast
dev1					
dev2					
dev3					
dev4					
dev5					
dev6					
dev7					

"Multicast" is Operating as Expected

	Topology	VLANs	PTP	BGP-Peer	Multicast
dev1					
dev2					
dev3					
dev4					
dev5					
dev6					
dev7					

"The Network" is Operating as Expected

	Topology	VLANs	PTP	BGP-Peer	Multicast
dev1					
dev2					
dev3					
dev4					
dev5					
dev6					
dev7					

Example Design Reports

Does the Design match the Requirements?

Design Reports:

- Reports ensure that Network Engineers can validate the network requirements at design time before proceeding with network configurations and deployments.
- Create cabling maps for installation teams
- Produces equipment lists of inventory assessments/ordering and pricing estimations
- Show how each port is being used within each design service context

Cabling Report

Device	Interface	Profile	Port	Remote Port	Remote Profile	Remote Interface	Remote Device
lf01.swb.aaa	Ethernet17	Leaf01SpineLagMember	SFPP-10G-LR	SFPP-10G-LR	SpineLeaf01LagMember	Ethernet43	sp01.swb.aaa
lf01.swb.aaa	Ethernet18	Leaf01SpineLagMember	SFPP-10G-LR	SFPP-10G-LR	SpineLeaf01LagMember	Ethernet44	sp01.swb.aaa
lf01.swb.aaa	Port-Channel2000	Leaf01SpineLag	virtual	virtual	SpineLeaf01Lag	Port-Channel2	sp01.swb.aaa
lf02.swb.aaa	Ethernet17	Leaf02SpineLagMember	SFPP-10G-LR	SFPP-10G-LR	SpineLeaf02LagMember	Ethernet45	sp01.swb.aaa
lf02.swb.aaa	Ethernet18	Leaf02SpineLagMember	SFPP-10G-LR	SFPP-10G-LR	SpineLeaf02LagMember	Ethernet46	sp01.swb.aaa
lf02.swb.aaa	Port-Channel2000	Leaf02SpineLag	virtual	virtual	SpineLeaf02Lag	Port-Channel3	sp01.swb.aaa
sp01.swb.aaa	Ethernet41	SpineStringerLagMember	SFPP-10G-LR	SFPP-10G-LR	StringerSpineLagMember	Ethernet17	string01.swb.aaa
sp01.swb.aaa	Ethernet42	SpineStringerLagMember	SFPP-10G-LR	SFPP-10G-LR	StringerSpineLagMember	Ethernet18	string01.swb.aaa
sp01.swb.aaa	Ethernet47	SpineTransitUplink	SFP-1G-T	CAT6-1G-RJ45	TransitSpineUplink	9	tr01.swb.aaa
sp01.swb.aaa	Ethernet48	SpineTransitUplink	SFP-1G-T	CAT6-1G-RJ45	TransitSpineUplink	10	tr01.swb.aaa
sp01.swb.aaa	Port-Channel1	SpineStringerLag	virtual	virtual	StringerSpineLag	Port-Channel2000	string01.swb.aaa

Interfaces

Interfaces assigned a ***Profile*** - identifies various design-service features and parameters

Examples include:

- VLANs used
- IP Addresses used
- PTP port parameters
- Physical port definition

sp01.swb.aaa: 55 interfaces

Name	Description	Profile	Port	Speed (Kbps)
Ethernet1	BDS-CAM01	BDSCamera	SFPP-10G-LR	10,000
Ethernet2	BDS-CAM02	BDSCamera	SFPP-10G-LR	10,000
Ethernet3	BDS-CAM03	BDSCamera	SFPP-10G-LR	10,000
Ethernet4	BDS-CAM04	BDSCamera	SFPP-10G-LR	10,000
Ethernet5	BDS-CAM05	BDSCamera	SFPP-10G-LR	10,000
Ethernet6	BDS-CAM06	BDSCamera	SFPP-10G-LR	10,000
Ethernet7	BDS-CAM07	BDSCamera	SFPP-10G-LR	10,000
Ethernet8	BDS-CAM08	BDSCamera	SFPP-10G-LR	10,000
Ethernet9	HET-PProcSRV01-HEI	HETPProcServerHEI	SFP28-25GBASE-CR	10,000
Ethernet10	HET-PProcSRV01-DATA1-1	HETPProcServerDATA	SFP28-25GBASE-CR	10,000
Ethernet11	HET-PProcSRV01-DATA2-1	HETPProcServerDATA	SFP28-25GBASE-CR	10,000
Ethernet12	HET-CtrlSRV02-HEI	HETCtrlServerHEI	SFP28-25GBASE-CR	10,000
Ethernet13	HET-CamSRV03-HEI	HETCamServerHEI	SFP28-25GBASE-CR	10,000
Ethernet14	HET-CamSRV03-DATA1-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet15	HET-CamSRV03-DATA2-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet16	HET-CamSRV04-HEI	HETCamServerHEI	SFP28-25GBASE-CR	10,000
Ethernet17	HET-CamSRV04-DATA1-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet18	HET-CamSRV04-DATA2-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet19	HET-CamSRV05-HEI	HETCamServerHEI	SFP28-25GBASE-CR	10,000
Ethernet20	HET-CamSRV05-DATA1-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet21	HET-CamSRV05-DATA2-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet22	HET-CamSRV06-HEI	HETCamServerHEI	SFP28-25GBASE-CR	10,000
Ethernet23	HET-CamSRV06-DATA1-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet24	HET-CamSRV06-DATA2-1	HETCamServerDATA	SFP28-25GBASE-CR	10,000
Ethernet30	node01.c01.k3s.swb.systems.mlbinfra.net-DAT1	SysinfraServerDATA	SFP28-25GBASE-CR	25,000
Ethernet31	node01.c01.k3s.swb.systems.mlbinfra.net-DAT2	SysinfraServerDATA	SFP28-25GBASE-CR	25,000
Ethernet32	node02.c01.k3s.swb.systems.mlbinfra.net-DAT1	SysinfraServerDATA	SFP28-25GBASE-CR	25,000
Ethernet33	node02.c01.k3s.swb.systems.mlbinfra.net-DAT2	SysinfraServerDATA	SFP28-25GBASE-CR	25,000
Ethernet41	string01.swb.aaa-et17	SpineStringerLagMember	SFPP-10G-LR	10,000
Ethernet42	string01.swb.aaa-et18	SpineStringerLagMember	SFPP-10G-LR	10,000
Ethernet43	lf01.swb.aaa-et17	SpineLeaf01LagMember	SFPP-10G-LR	10,000
Ethernet44	lf01.swb.aaa-et18	SpineLeaf01LagMember	SFPP-10G-LR	10,000
Ethernet45	lf02.swb.aaa-et17	SpineLeaf02LagMember	SFPP-10G-LR	10,000

PTP Port Usage Example

PTP domain 0
41 ports from 4 devices

Port State(s)	Device	Interface	Desc	Kind
Grandmaster | Slave | Master	string01.swb.aaa	Ethernet8	BLANT01	BoleroAntenna
Grandmaster | Slave | Master	string01.swb.aaa	Ethernet9	BLANT02	BoleroAntenna
Slave | Passive	lf02.swb.aaa	Port-Channel2000	sp01.swb.aaa-po3	Leaf02SpineLag
Slave | Passive	sp01.swb.aaa	Port-Channel2	lf01.swb.aaa-po2000	SpineLeaf01Lag
Slave | Passive	sp01.swb.aaa	Port-Channel1	string01.swb.aaa-po2000	SpineStringerLag
Passive | Master	string01.swb.aaa	Port-Channel2000	sp01.swb.aaa-po1	StringerSpineLag
Passive | Master	sp01.swb.aaa	Port-Channel3	lf02.swb.aaa-po2000	SpineLeaf02Lag
Master	sp01.swb.aaa	Ethernet1	BDS-CAM01	BDSCamera
Master	sp01.swb.aaa	Ethernet2	BDS-CAM02	BDSCamera
Master	sp01.swb.aaa	Ethernet3	BDS-CAM03	BDSCamera
Master	sp01.swb.aaa	Ethernet4	BDS-CAM04	BDSCamera
Master	sp01.swb.aaa	Ethernet5	BDS-CAM05	BDSCamera
Master	sp01.swb.aaa	Ethernet6	BDS-CAM06	BDSCamera
Master	sp01.swb.aaa	Ethernet7	BDS-CAM07	BDSCamera
Master	sp01.swb.aaa	Ethernet8	BDS-CAM08	BDSCamera
Master	sp01.swb.aaa	Ethernet9	HET-PProcSRV01-HEI	HETPProcServerHEI
Master	sp01.swb.aaa	Ethernet10	HET-PProcSRV01-DATA1-1	HETPProcServerDATA
Master	sp01.swb.aaa	Ethernet11	HET-PProcSRV01-DATA2-1	HETPProcServerDATA
Master	sp01.swb.aaa	Ethernet12	HET-CtrlSRV02-HEI	HETCtrlServerHEI

Design Services

Reusable building blocks

Design Services - Reusable Building Blocks

- Network Vendor/System agnostic
- Defines design ***validation rules*** so that we cannot build "invalid" designs
- Provides service specific ***build automations***
- Defines ***collections of checks*** that will be generated for each device using the design service
- Can be general purpose such as "topology" and "vlans"
- Can be bespoke to enterprise specific needs

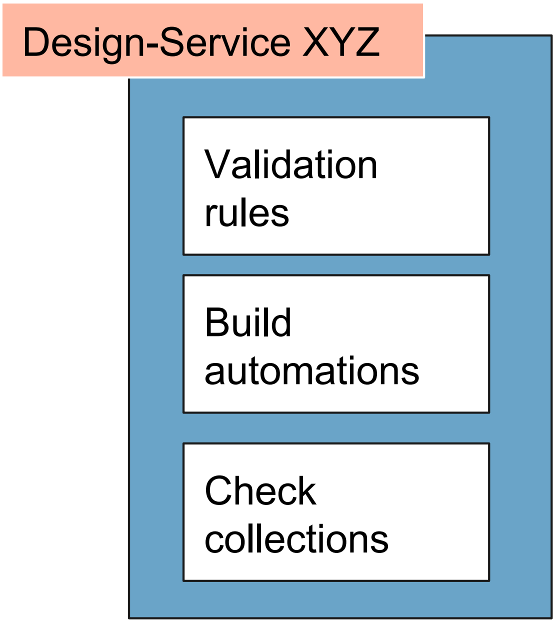

Design Automation - VLAN Port Assignments

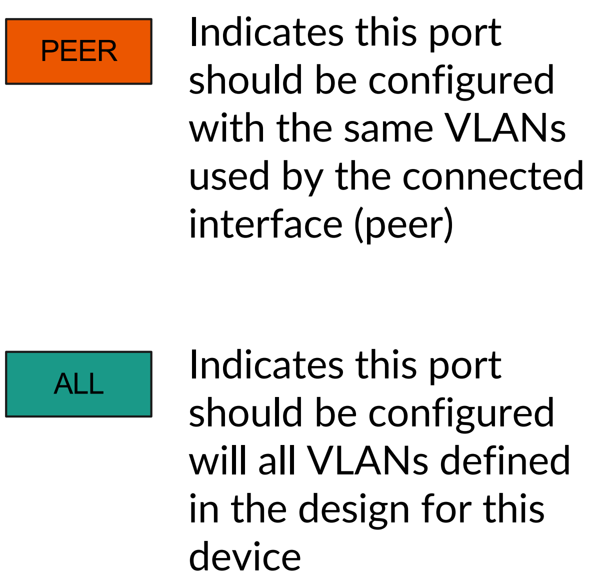

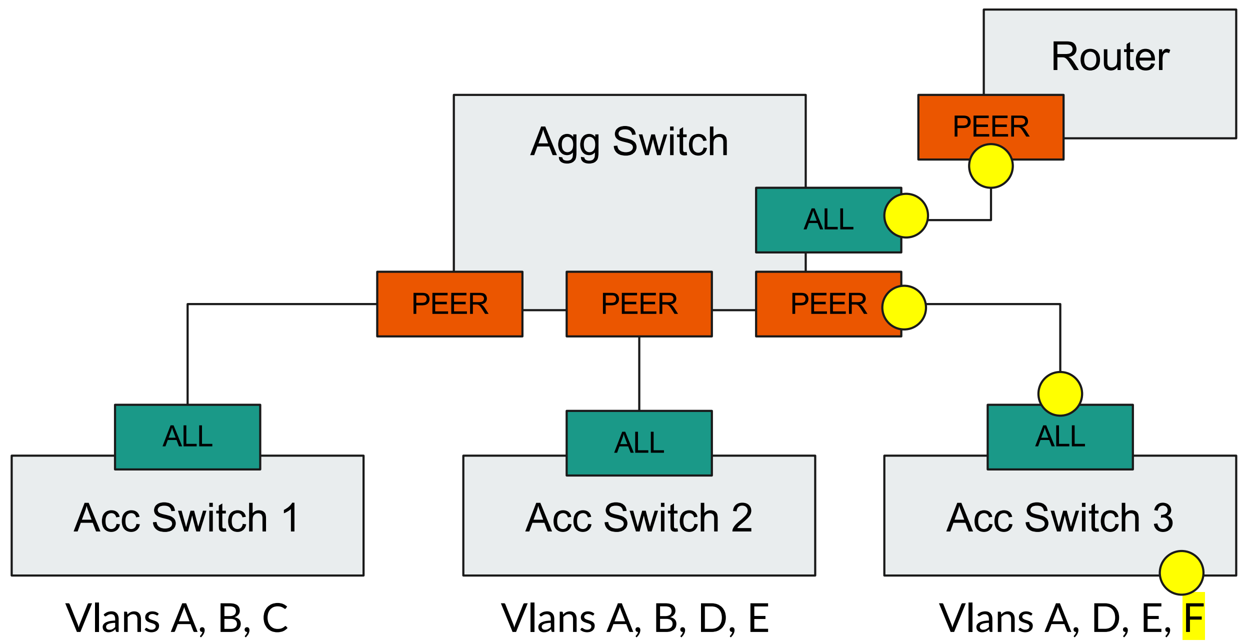

Other Design Validation Examples

- Ensure both ends of a cable are "connected" somewhere
- Ensure both sides of a cable are of the same form and speed
- Ensure that transceivers are used properly
  - Designed into interfaces that have ports that require them
  - Or vice-versa
- Ensure both sides of a BGP point-to-point session are in the same subnet

Network Validations

Is your network **operating state** running as expected?

- Validations by Design
- Physical Topology
- Device Services
- Network Wide

Anatomy of a Validation

The Check/Measurement definition is a unit of abstraction

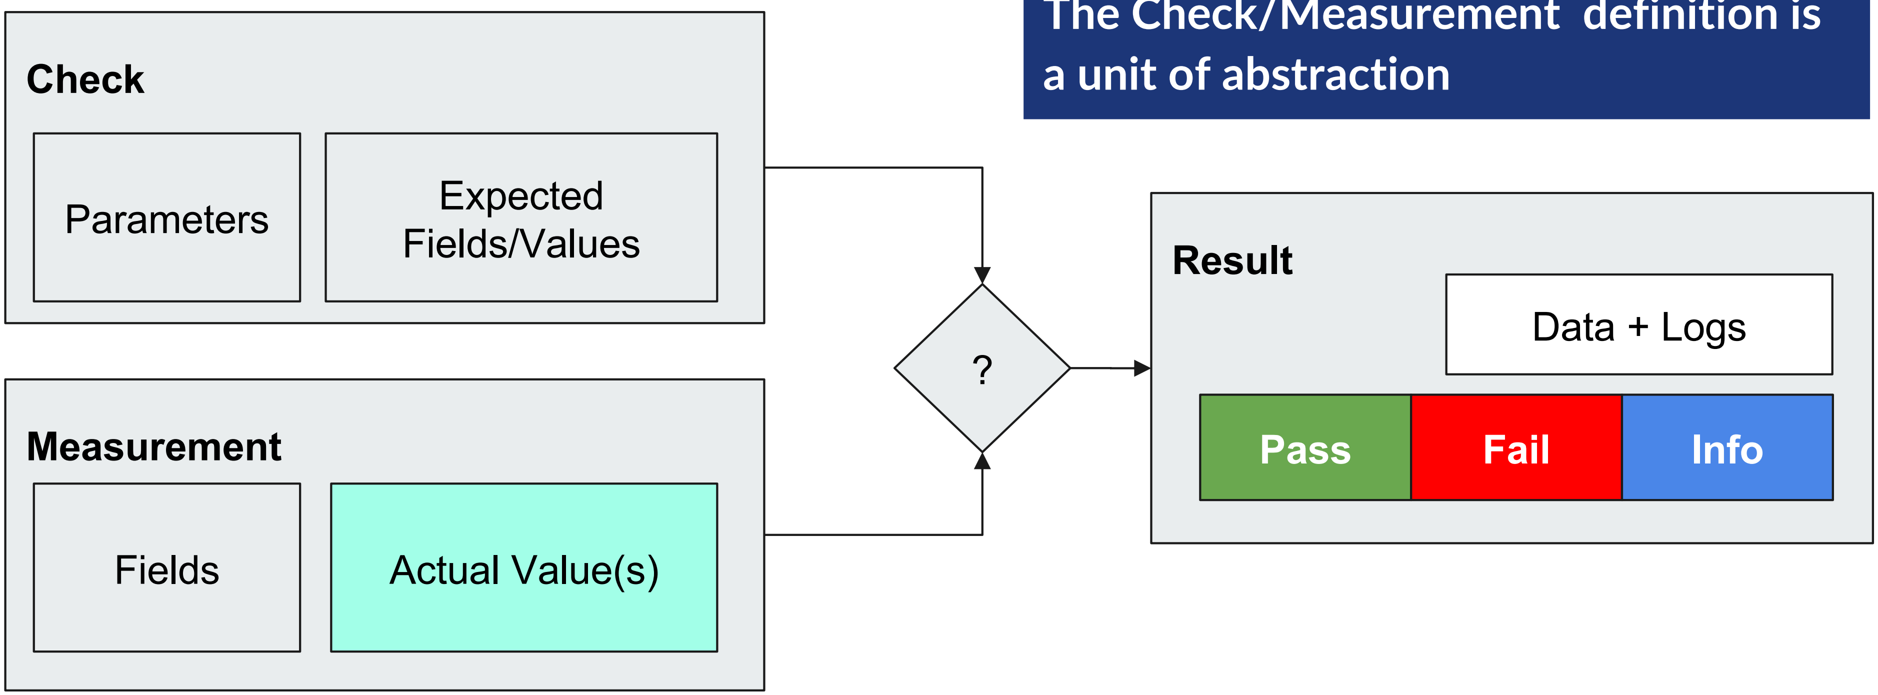

Validations By Design-Service

Checks are **automatically** generated based on the design

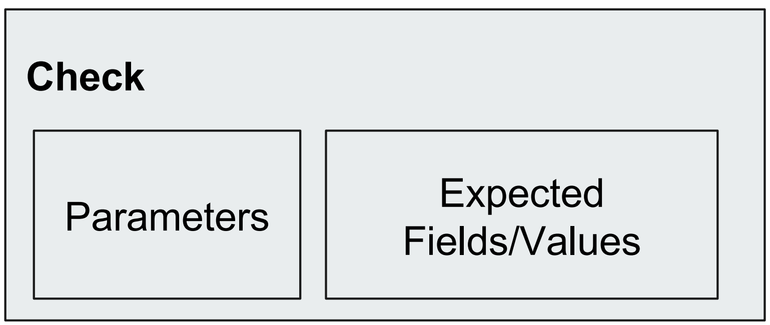

- Each design-service defines their check collections
- When the design changes the collection of checks change
- Parameters change
- Expected results change
- Add / remove checks

"Measurement Engines" Implements Service Checks

A measurement engine represents the per-Vendor/OS + per Design-Service Check level of abstraction

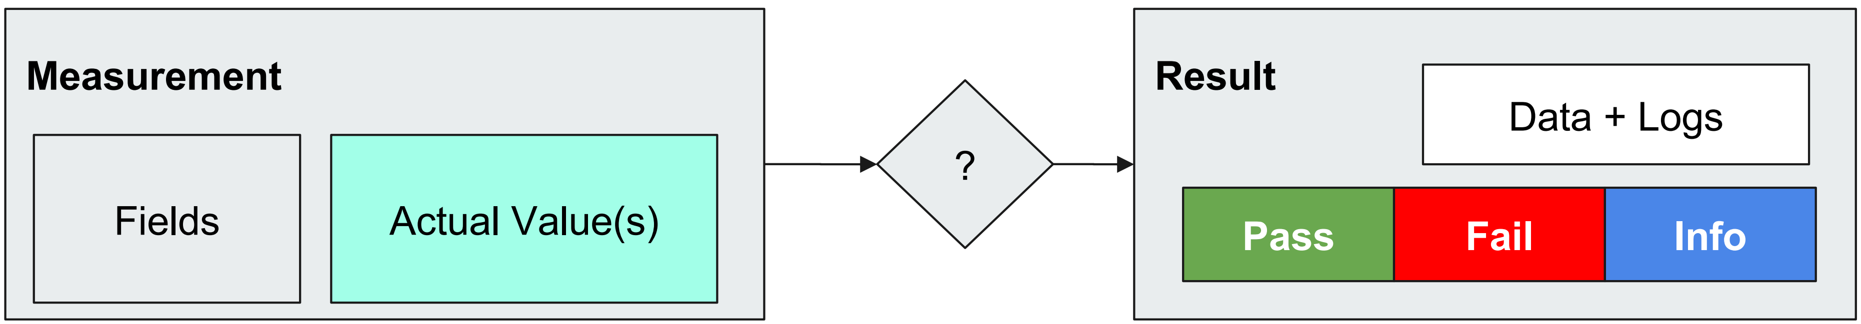

IOS, IOS-XE, EOS, Junos, PyATS, SuzieQ, etc ….

Design Service Validation Abstraction Example

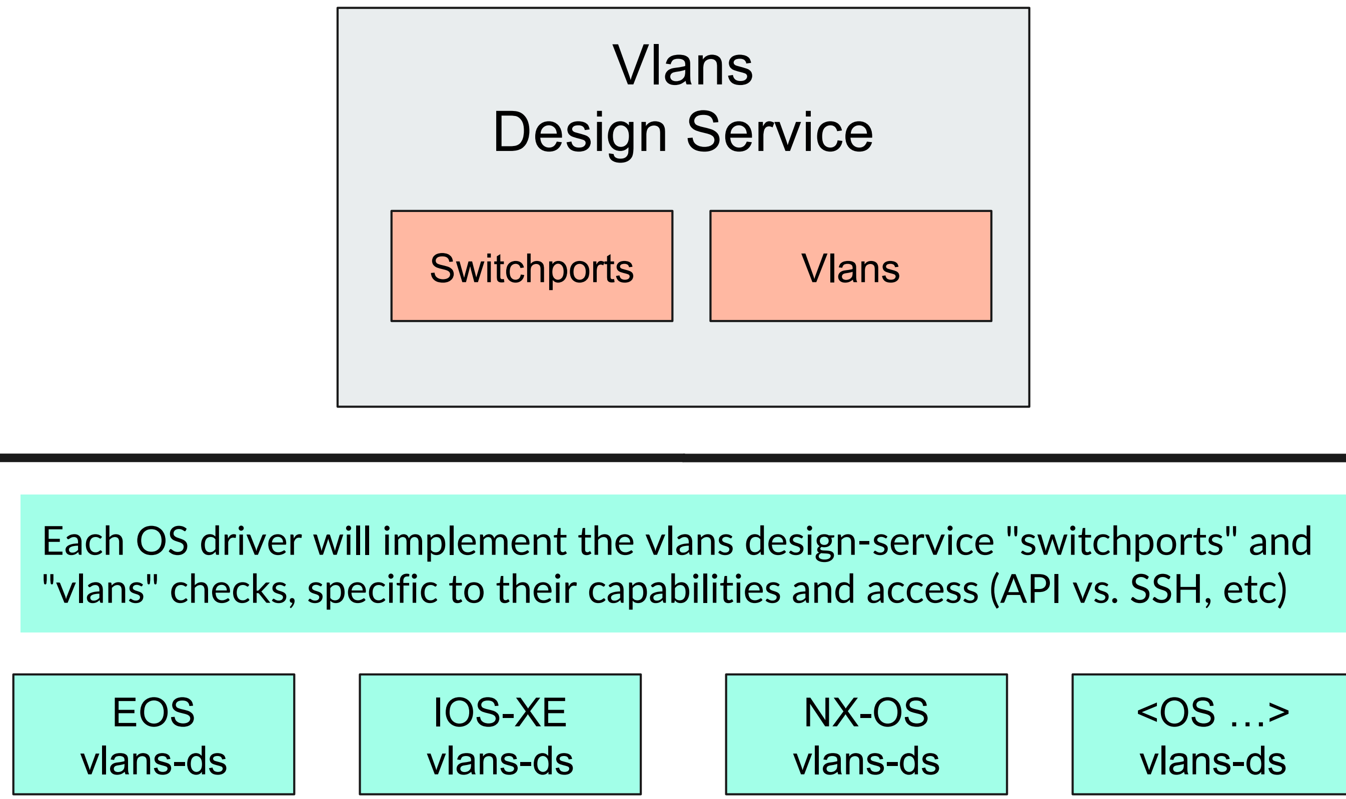

Topology Design Service checks ...

- Device is reachable ... via API
- Device is the model / platform as expected
- Interfaces cabled as expected
- Interfaces using the expected transceivers
- Interface up that should be up
- Interfaces down that should be down
- Interfaces operating speed as expected (1g vs 10g)

Example - Interface Checks Failing

```
Device: lf01.lhv.aaa
   Test Logs: interfaces.json
```

Status	Device	Id	Field	Log
FAIL	lf01.lhv.aaa	Ethernet3	oper_up	FAIL oper_up {'expected': True, 'measured': False} SKIP speed {'expected': 1000, 'measured': 0} PASS used True PASS desc 'BBR01'
FAIL	lf01.lhv.aaa	Ethernet4	speed	FAIL speed {'expected': 1000, 'measured': 10} PASS used True PASS desc 'HET-CtrlSRV02-MGMT' PASS oper_up True

Example - Interface Cabling Passing

```
Device: lf01.lhv.aaa
   Test Logs: cabling.json
```

Status	Device	Id	Field	Log
PASS	lf01.lhv.aaa	Ethernet1	device	PASS device 'string01.lhv.aaa' PASS port_id 'Ethernet17'
PASS	lf01.lhv.aaa	Ethernet2	device	PASS device 'string01.lhv.aaa' PASS port_id 'Ethernet18'
PASS	lf01.lhv.aaa	Ethernet15	device, port_id	PASS device 'sp01.lhv.aaa' PASS port_id 'GigabitEthernet1/0/21'
PASS	lf01.lhv.aaa	Ethernet16	device, port_id	PASS device 'sp01.lhv.aaa' PASS port_id 'GigabitEthernet1/0/22'

Example - Collecting Informational Measurements

```
Device: lf01.tol.aaa
  Test Logs: device.json
```

Status	Device	Id	Field	Log
INFO	lf01.tol.aaa	lf01.tol.aaa		mfgName 'Arista' modelName 'CCS-710P-16P' hardwareRevision '11.01' serialNumber [redacted] systemMacAddress [redacted] hwMacAddress [redacted] configMacAddress [redacted] version '4.29.1F' architecture 'i686' internalVersion '4.29.1F-29834883.4291F' internalBuildId '025c9e79-f2f5-4361-81ce-5fdfc4278b54' imageFormatVersion '3.0' imageOptimization 'Strata-4GB' bootupTimestamp 1679661446.8029 uptime 366513.29 memTotal 3982416 memFree 2581316 isIntlVersion False

Device Network Quality Assurance

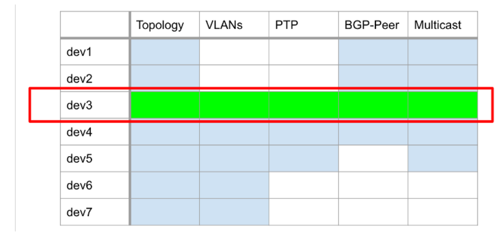

- Are VLANs on switch-ports as expected
- Are PTP port-states in the expected state
- Is there one and only one interface accepting PTP clocking
- Is the PTP clocking class the correct value
- Are all expected BGP neighbors up as expected
- Are all multicast IP sources from source-hosts as expected

	Topology	VLANs	PTP	BGP-Peer	Multicast
dev1					
dev2					
dev3					
dev4					
dev5					
dev6					
dev7					

Network Wide Quality Assurance

- Are all devices in the same PTP domain reporting the same PTP grandmaster ID
- Are PTP port-states correct by the rules of PTP
- Are LAG pairs reporting consistent bundling between devices
- Are routes showing up as expected across the network
- Are multicast S,G paths sources-receivers across the network as expected

Example PTP Network Wide Design Report

The design describes the expected PTP state(s) for each device and interface

Port State(s)	Device	Interface	Desc	Kind
Grandmaster | Slave | Master	string01.lhv.aaa	Ethernet8	BLANT01	[illegible]
Grandmaster | Slave | Master	string01.lhv.aaa	Ethernet9	BLANT02	[illegible]
Slave | Passive	sp01.lhv.aaa	Port-channel2	lf01.lhv.aaa-po2000	[illegible]
Slave | Passive	lf01.lhv.aaa	Port-Channel1	string01.lhv.aaa-po2000	[illegible]
Passive | Master	string01.lhv.aaa	Port-Channel2000	lf01.lhv.aaa-po1	[illegible]
Passive | Master	lf01.lhv.aaa	Port-Channel2000	sp01.lhv.aaa-po2	[illegible]
Master	lf01.lhv.aaa	Ethernet3	BBR01	[illegible]
Master	lf01.lhv.aaa	Ethernet4	HET-CtrlSRV02-MGMT	[illegible]
Master	lf01.lhv.aaa	Ethernet5	HET-CamSRV03-MGMT	[illegible]
Master	lf01.lhv.aaa	Ethernet6	HET-CamSRV04-MGMT	[illegible]
Master	lf01.lhv.aaa	Ethernet7	HET-CamSRV05-MGMT	[illegible]
Master	lf01.lhv.aaa	Ethernet8	HET-CamSRV06-MGMT	[illegible]
Master	lf01.lhv.aaa	[illegible]	[illegible]	[illegible]
Master	lf01.lhv.aaa	[illegible]	[illegible]	[illegible]
Master	lf01.lhv.aaa	[illegible]	[illegible]	[illegible]
Master	lf01.lhv.aaa	Ethernet17	ubiquity1-port8	[illegible]
Master	lf01.lhv.aaa	Ethernet18	ubiquity2-port8	[illegible]

What is the actual state of PTP in this network ?

Example PTP Network Wide Report

Site where PTP service is **NOT** working as expected.
Report shows informational and potential corrective actions.

Service: ptp, 5 logs

Status	Title	Message
PASS	Grandmaster ID	Consistent grandmaster ID 2c:dd:e9:ff:ff:fd:6b:08 across devices
PASS	PTP port peering checks	All port-to-port peering checks PASS
INFO	Port disabled-down	lf01.lhv.aaa, Ethernet3 port state is disabled becauase interface is down
INFO	Port disabled-down	string01.lhv.aaa, Ethernet9 port state is disabled becauase interface is down
FAIL	Grandmaster Port	Gradmaster device not found. Expected on: - string01.lhv.aaa: Ethernet8 - string01.lhv.aaa: Ethernet9

Example PTP Network Wide Report

Site where PTP service is working as expected

Service: ptp, 4 logs

Status	Title	Message
PASS	Grandmaster Port	Grandmaster device found connected on string01.drm.aaa, Ethernet8
PASS	Grandmaster ID	Consistent grandmaster ID 00:19:7c:ff:fe:09:8f:35 across devices
PASS	PTP port peering checks	All port-to-port peering checks PASS
PASS	Any PTP checks fail	All check results PASS.

Design Driven Network Assurance

In Summary …

- Design is "code" - version controlled, reviewed, lifecycle
- Design checks are quality assurance tests that provide operational coverage
- Encourages behavior to keep Designs "up to date" by the nature of running the network by Design
- Builds skill sets to manage IaC based systems

Q & A

Thank you!