
cisco Live!
ALL IN

The bridge to possible

Simple Leaf/Spine with a Touch of ToR

Network Designs for the Modern Data Center

Brenden Buresh – Principal Architect
@BrendenBuresh

BRKDCN-2229

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

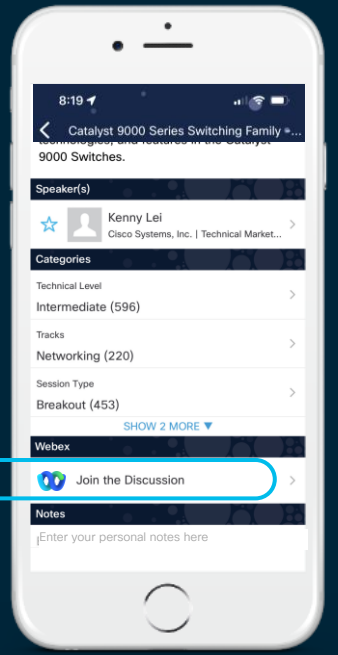

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2229

Webex spaces will be moderated by the speaker until June 17, 2022.

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- Why Did We Introduce FEX?
- The Evolution of DC Network Designs
- Bandwidth/Cost Evolution Over a Decade
- Migration Considerations
- Conclusion

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Why Did We Introduce FEX?

Middle of Row (MoR) and End of Row (EoR)

Big Centralized Chassis

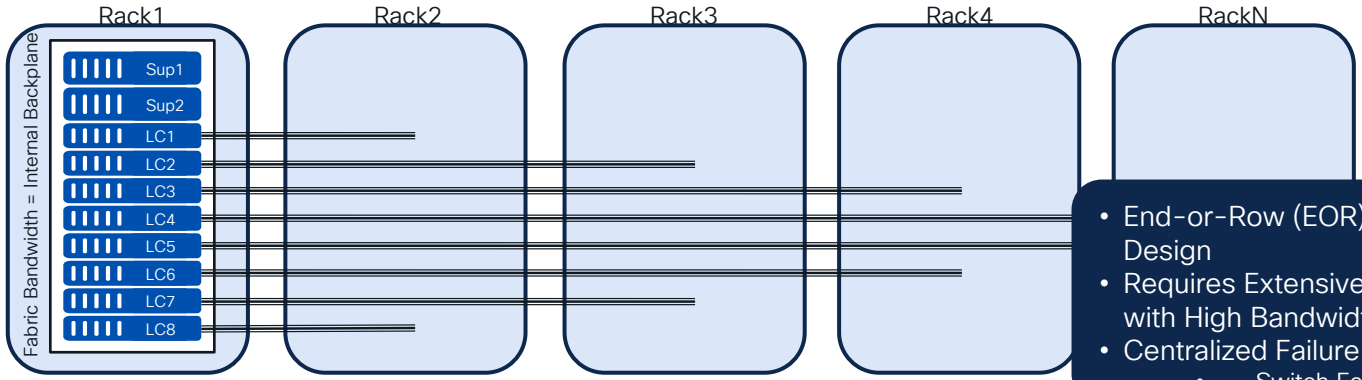

- End-or-Row (EOR) or Mid-of-Row (MOR) Design
- Requires Extensive Structured Ethernet Cabling with High Bandwidth Switch
- Centralized Failure Domain
  - Switch Failure = Network Failure

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

What is FEX?

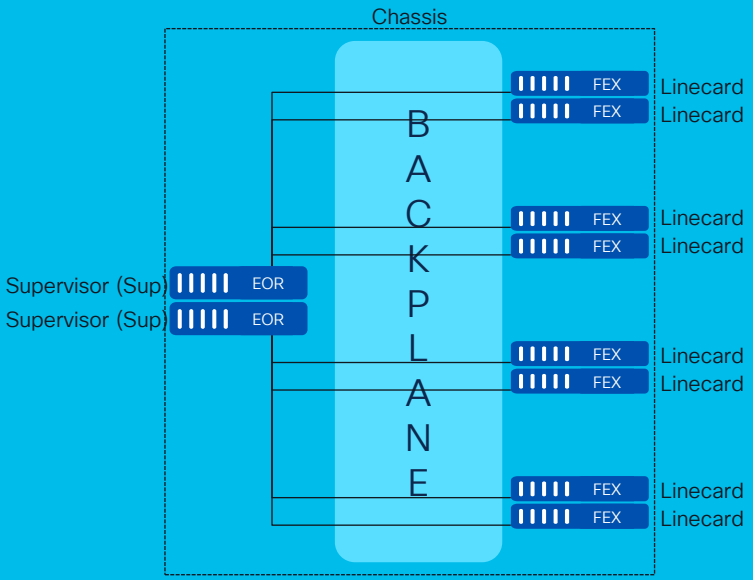

- A FEX can be seen as a way of “disaggregating” a traditional modular switch
- Enables the capability to build a centrally managed but highly distributed network design

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Middle of Row (MoR) and End of Row (EoR)

From Big Centralized to FEX

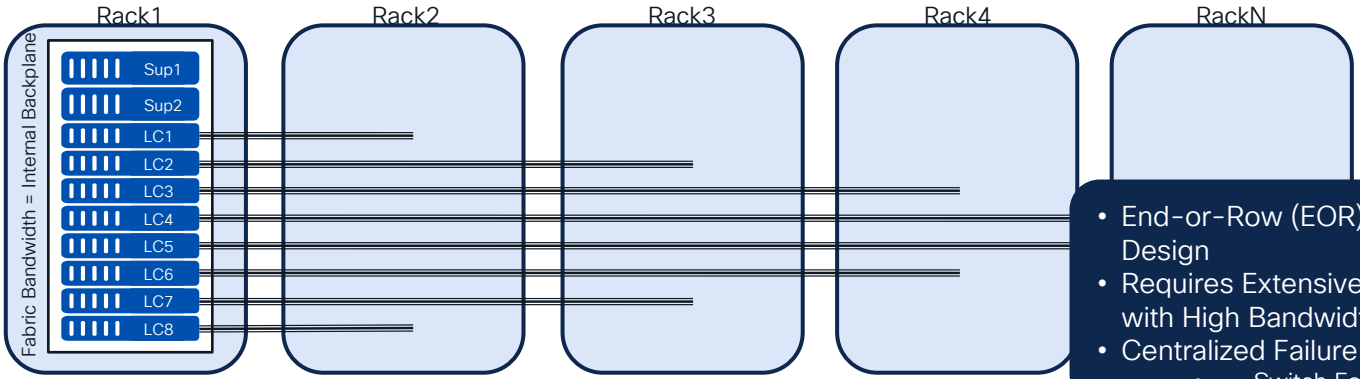

- End-or-Row (EOR) or Mid-of-Row (MOR) Design
- Requires Extensive Structured Ethernet Cabling with High Bandwidth Switch
- Centralized Failure Domain
  - Switch Failure = Network Failure

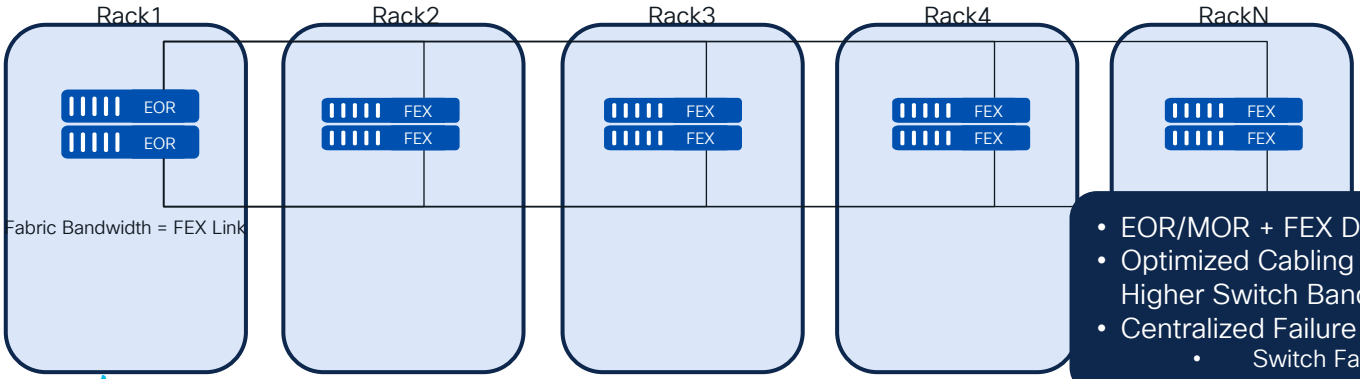

- EOR/MOR + FEX Design
- Optimized Cabling Requirement and Leverages Higher Switch Bandwidth
- Centralized Failure Domain
  - Switch Failure = Network Failure

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Why Did We Introduce FEX?

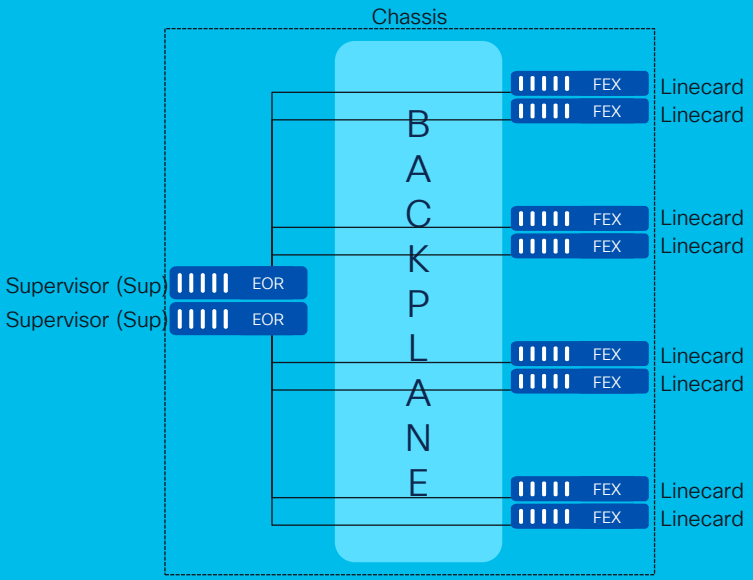

N5k: 24 FEX * 48 Host Ports = 1152 Host Ports (HIF)
N9k: 16 FEX * 48 Host Ports = 768 Host Ports (HIF)

- Centralized Management
- Modular Chassis Feeling
  - Unified CLI Structure for Config and Operation
- Capability of offering multiple port speeds (100M/1G/10G)
- Economics, Relative High Cost of Switch Ports or $ per Gbps

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

When to Avoid Leveraging FEX?

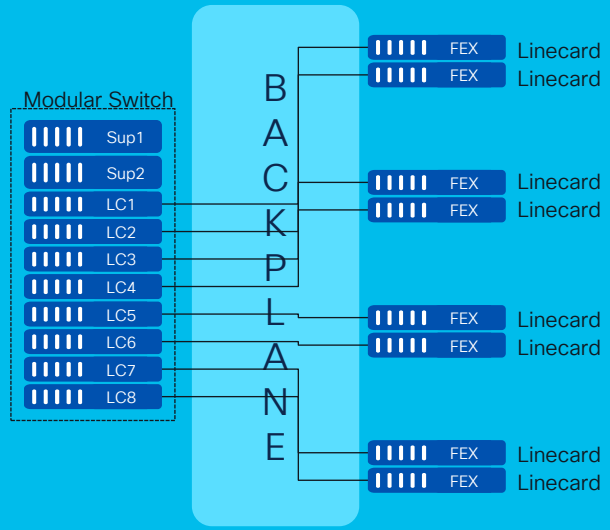

N7k: 64 FEX * 48 Host Ports = 3072 Host Ports (HIF)

- Extending Centralized Management (Beyond the Linecards)
- Increasing Modular Chassis Reach
  - Nested Linecard
  - Extending Failure Domain
- Giving Up the Benefits of a Distributed Fabric

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

A Data Center Fabric Prior to Data Center Fabrics

Cisco Fabric Extender (FEX) Overview

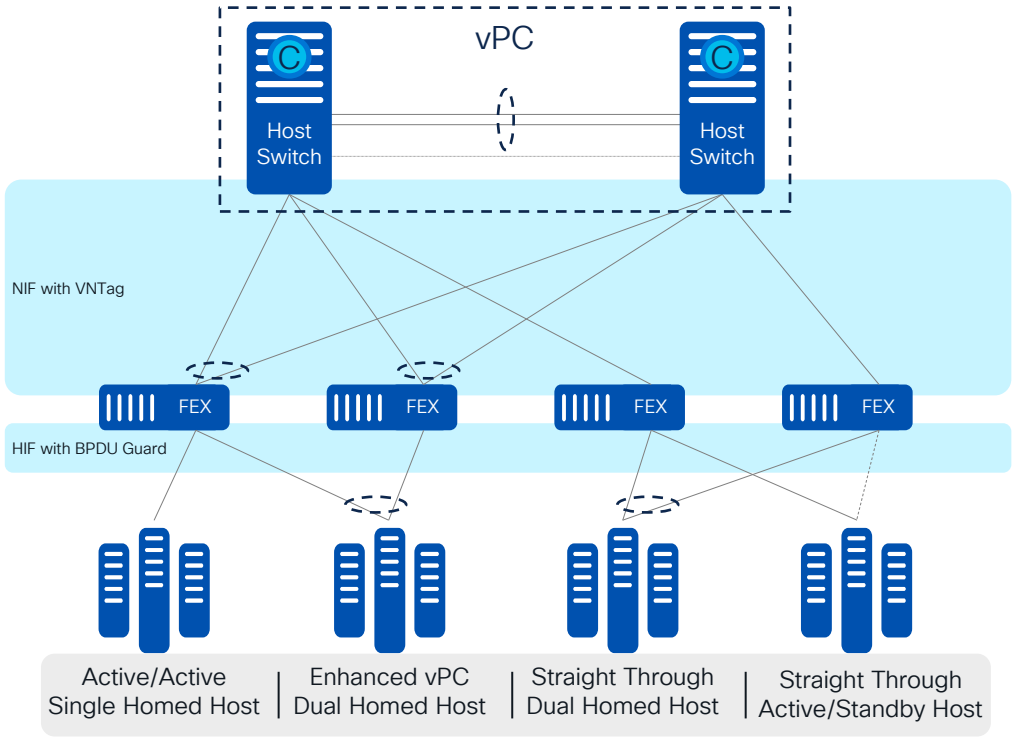

13 Years ago
Around 2009

- Centralized Management
  - Co-located on the Switch
  - Limited to No Synchronization
  - Host Switch Operational Dependency
- Network Redundancy (NIF to NIF)
  - Uses VNTag (802.1BR / 802.1Qbh)
  - 1+1 Redundancy based on Layer-2 Port-Channel (vPC)
- Host Redundancy (Host to HIF)
  - Single Homed or Dual Homed Hosts (vPC, A/S)
  - Spanning-Tree BPDU Guard
  - Subset of HIF Capabilities (Dependent on Host Switch)

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

The Evolution of DC Network Designs

A Data Center Fabric Prior to Data Center Fabrics

Cisco Fabric Extender (FEX) Overview

13 Years ago
Around 2009

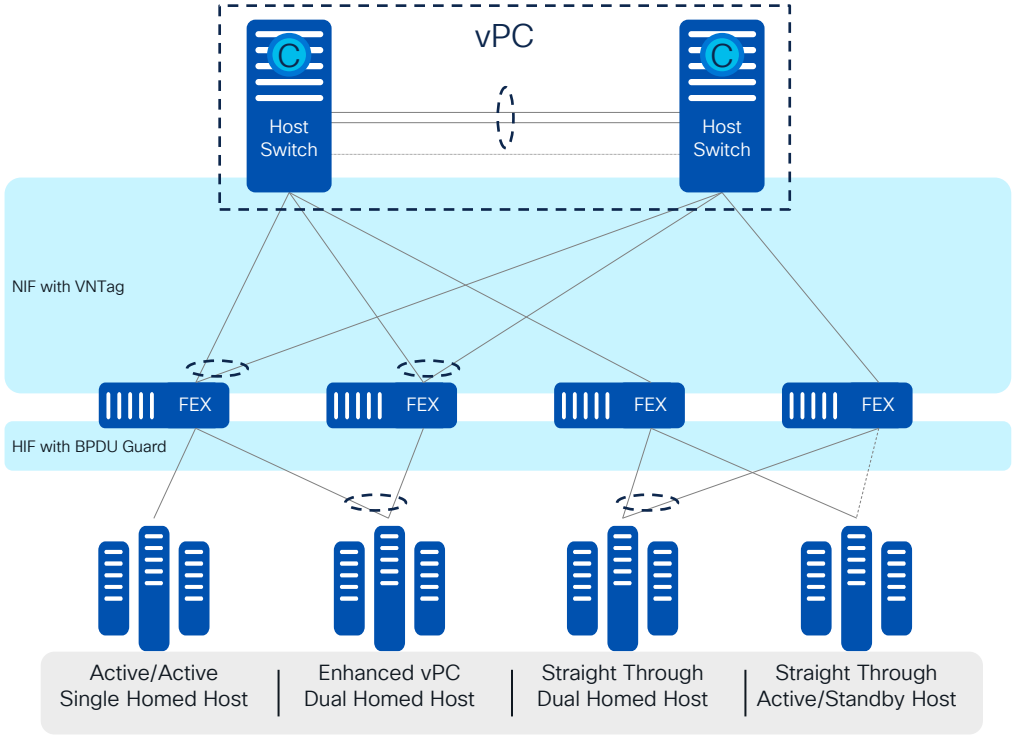

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

A Data Center Fabric Prior to Data Center Fabrics

Evolution to a Fabric

11 Years ago
Around 2011

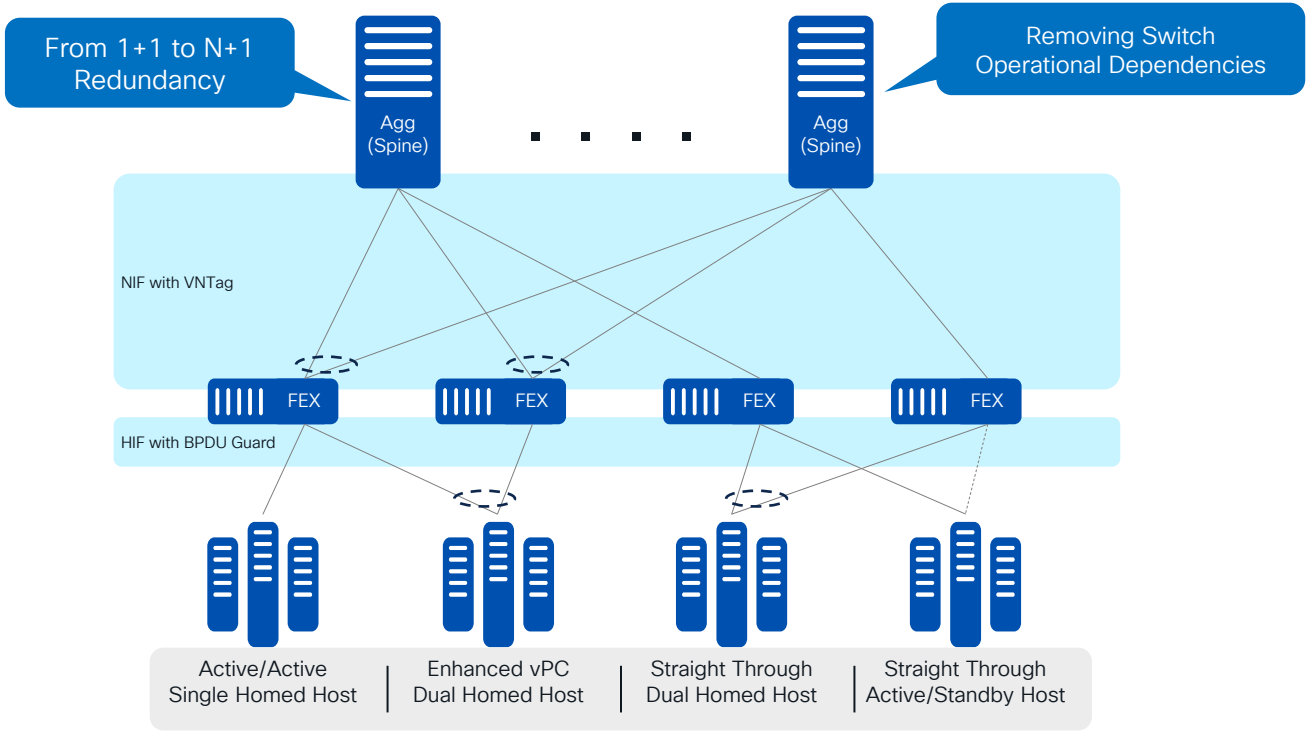

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Early Steps in the Data Center Fabric Evolution

Evolution to a Fabric

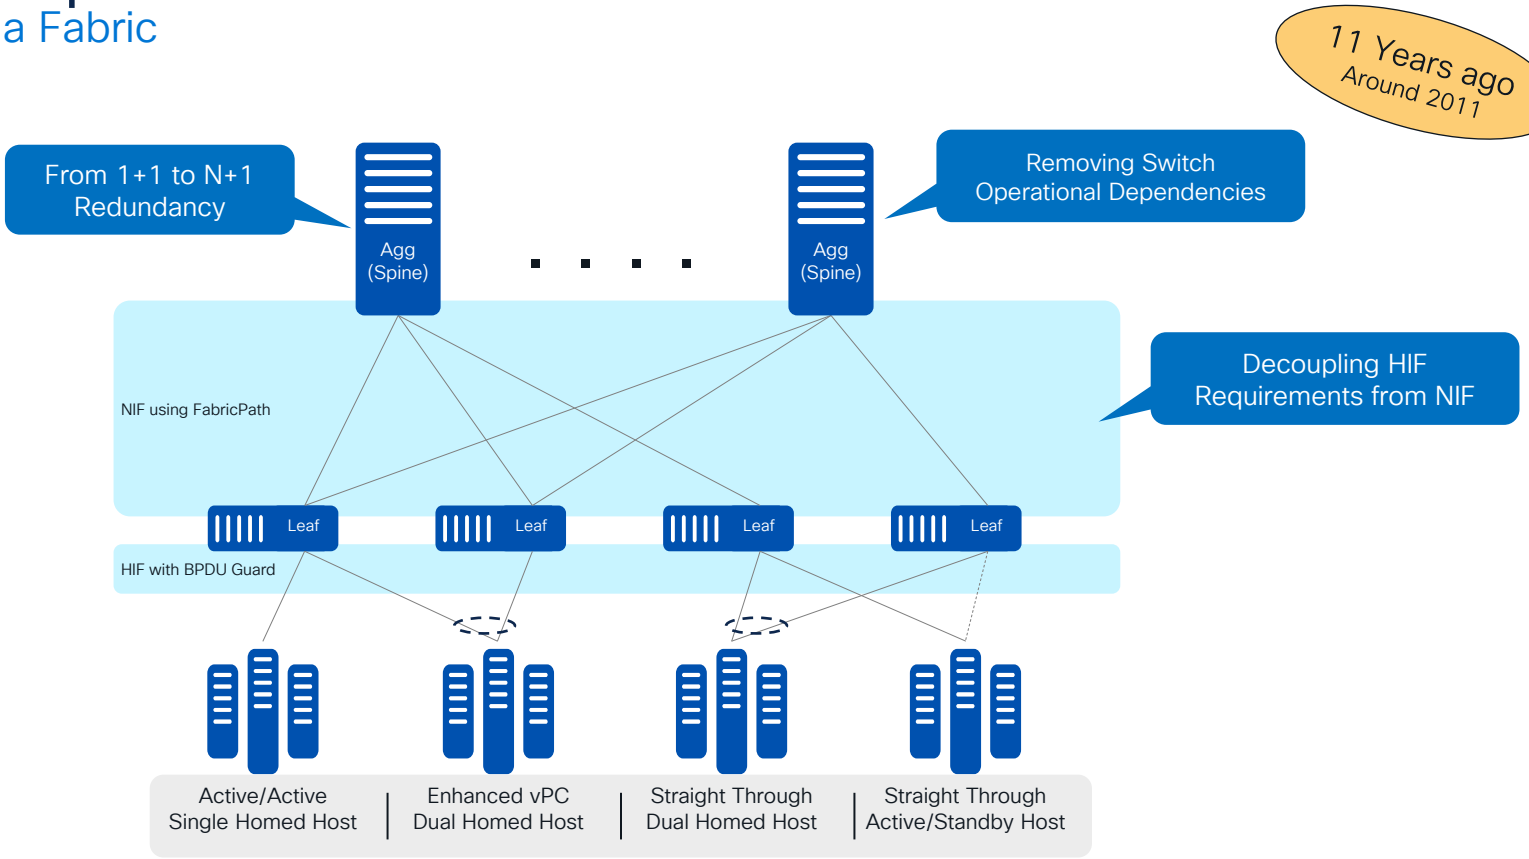

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Early Steps in the Data Center Fabric Evolution

Evolution to a Fabric

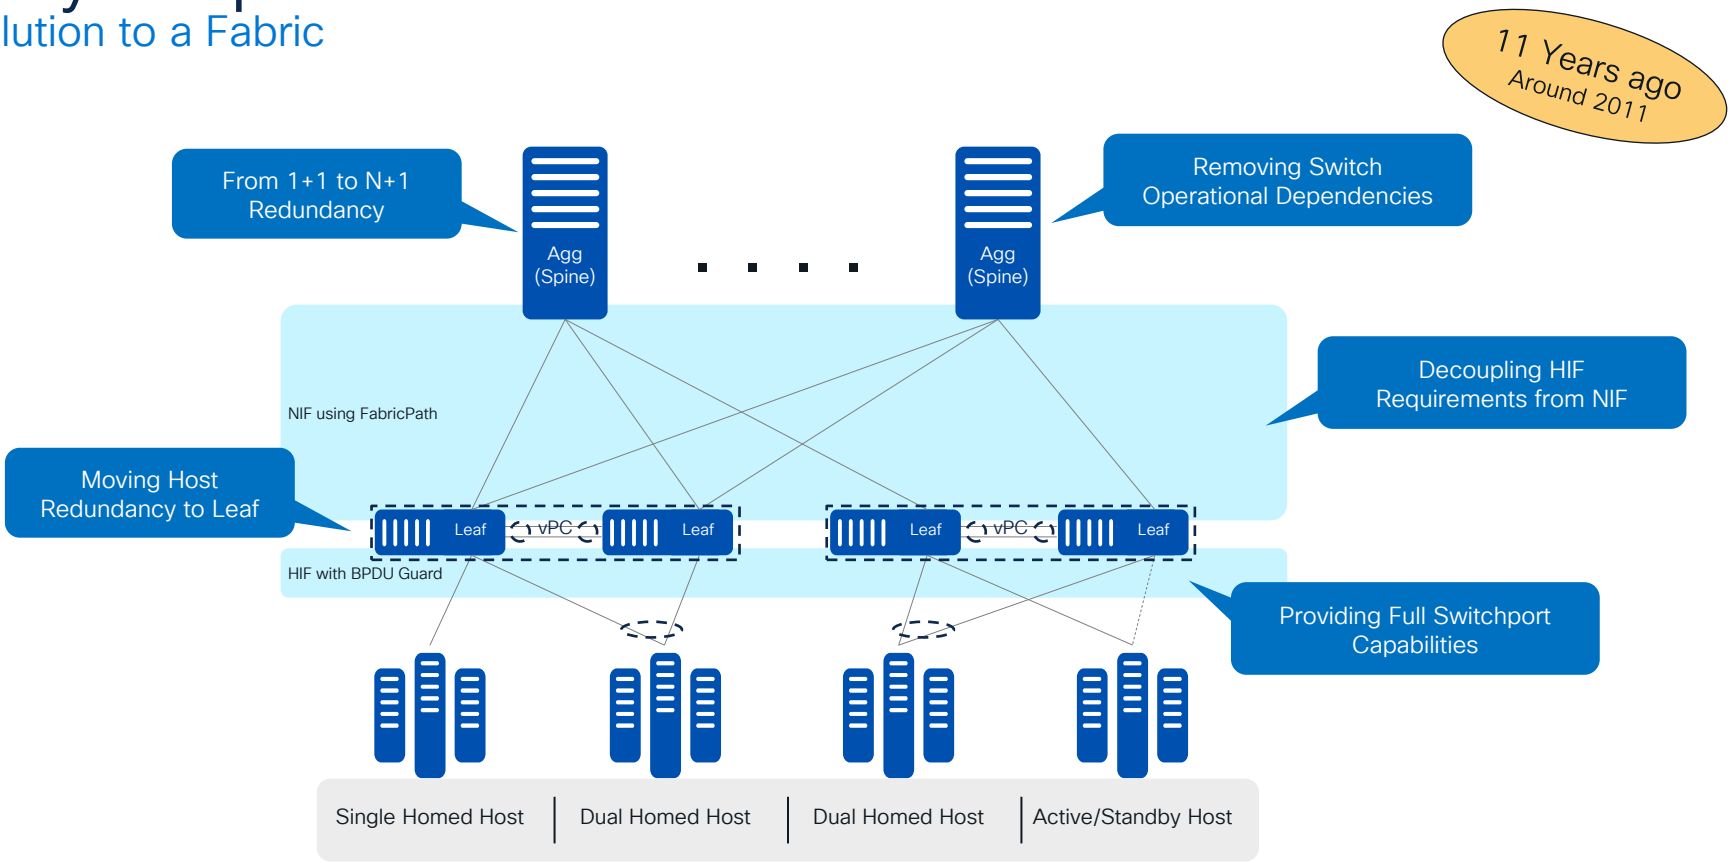

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Early Steps in the Data Center Fabric Evolution

Cisco FabricPath Overview

11 Years ago
Around 2011

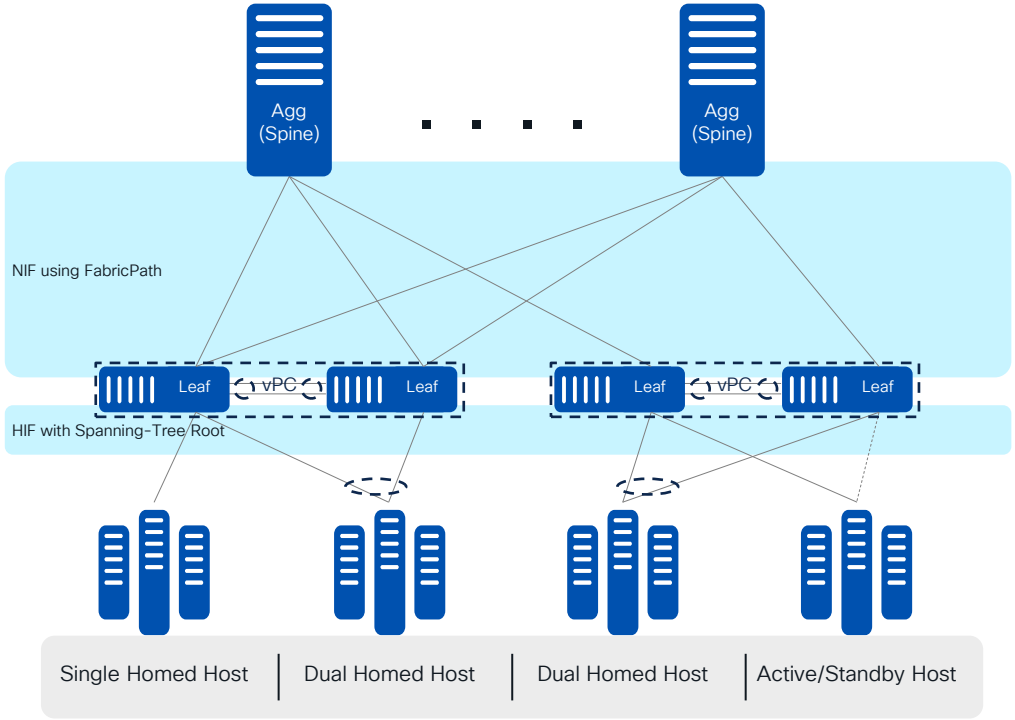

- Centralized Management
  - Nothing Really There
- Network Redundancy (Leaf to Spine)
  - FabricPath (MAC-in-MAC), requires Agg/Spine Support
  - N+1 Redundancy with ECMP
- Host Redundancy (Host to Leaf)
  - Single Homed or Dual Homed Hosts (vPC, A/S)
  - Full HIF Capabilities at Leaf with Spanning-Tree Root

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using Mature SDN for Data Center Fabrics

Evolution to Today's Fabric

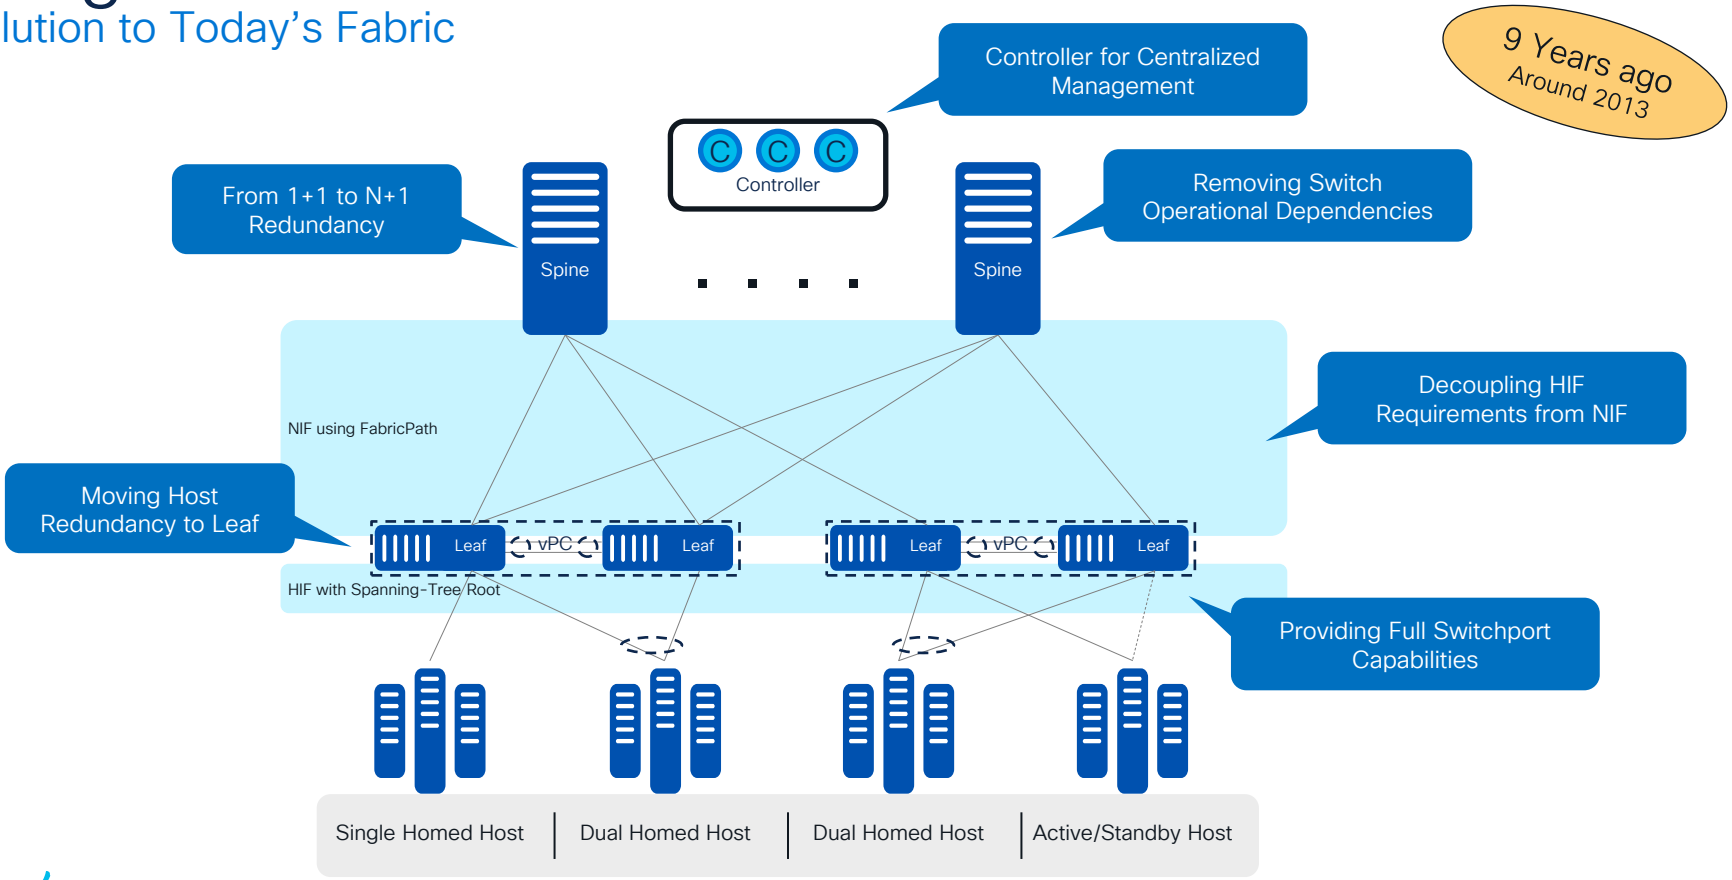

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using Mature SDN for Data Center Fabrics

Evolution to Today's Fabric

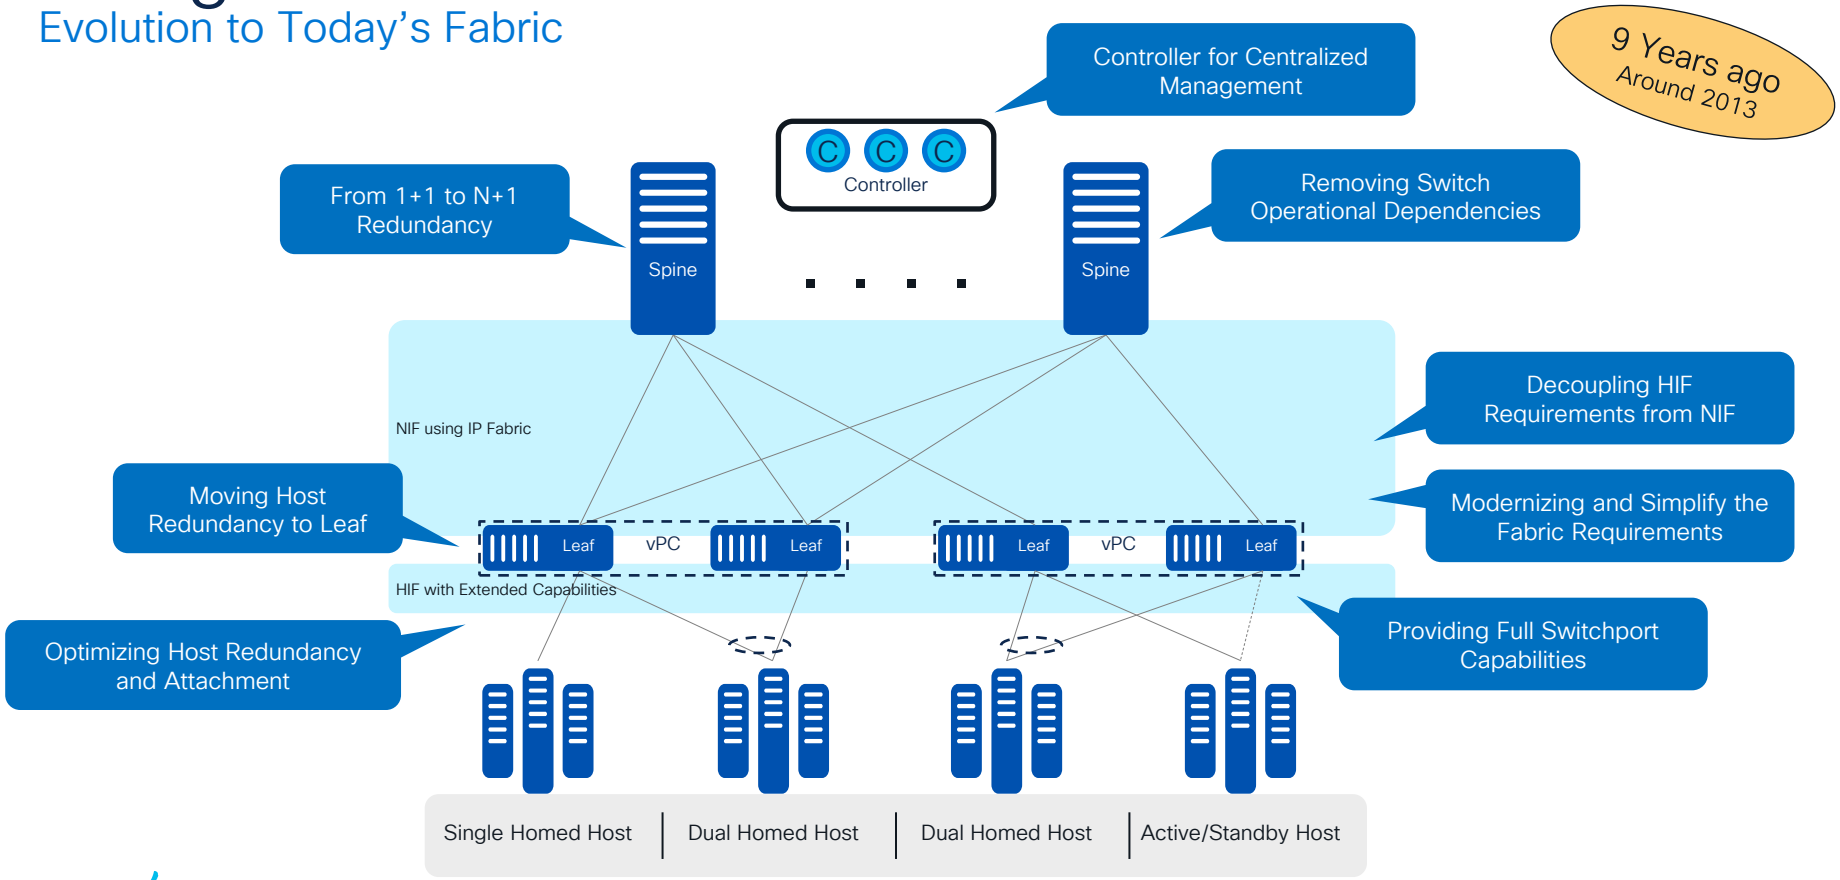

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Using Mature SDN for Data Center Fabrics

Cisco ACI and VXLAN EVPN Fabric Overview

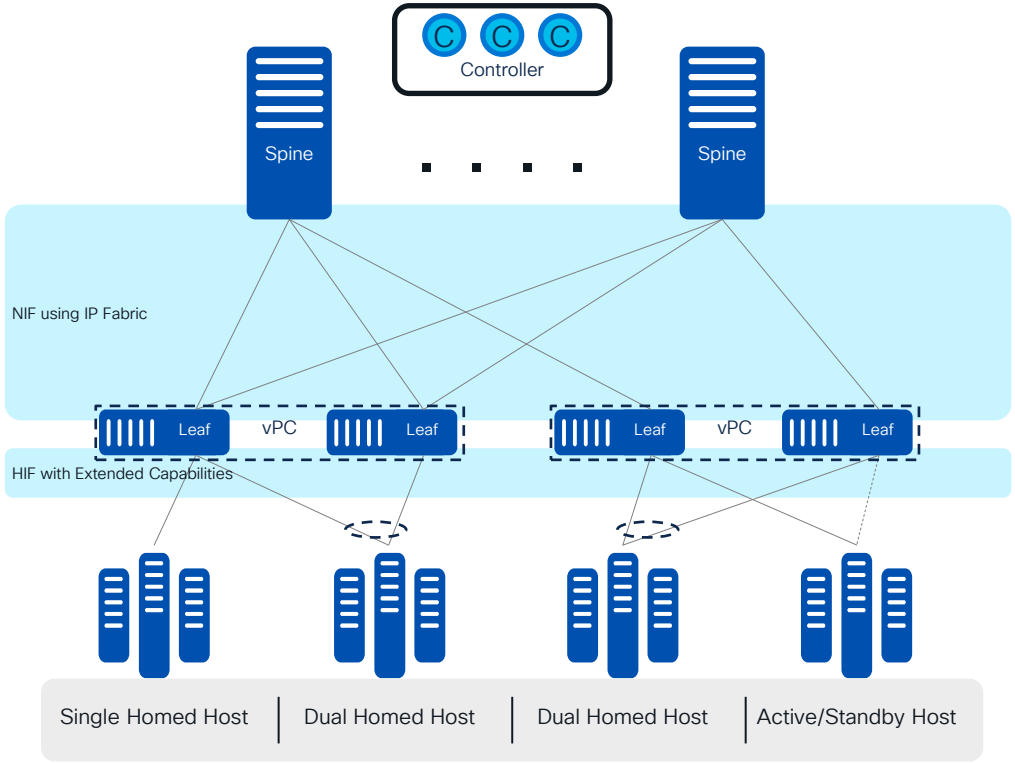

9 Years ago
Around 2013

- Centralized Management
  - Independent to Switch Operating System
  - Full Config Synchronization
  - N+1 Cluster or High-Availability
- Network Redundancy (Leaf to Spine)
  - Uses VXLAN (RFC7348), the Spine is just an IP Router
  - N+1 Redundancy based on IP Fabric (ECMP)
- Host Redundancy (Host to Leaf)
  - Single Homed or Dual Homed Hosts (vPC, A/S)
  - Full HIF Capabilities

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Bandwidth/Cost Evolution Over a Decade

Pricing Economics

Nexus 5000 + FEX

Before 2013

Host Switch	Price USD$	$ per Gbps
Cisco Nexus 5672UP-16G • 48 x 10G + 6 x 40G	USD $27,600.00	USD ~$38.00/Gbps
(48*10Gbps)+(6*40Gbps) =720Gbps/Switch	USD $27,600.00/720Gbps	

FEX	Price USD$	$ per Gbps
Cisco Nexus 2348UPQF-QSA • 48 x 1G + 6 x 10G	USD $24,300.00	USD $506.00/Gbps
48*1Gbps =48Gbps HIF per FEX	USD $24,300.00/48Gbps	

FEX Port $ per Gbps (1.2 Oversubscription)

USD ~$544.00/Gbps

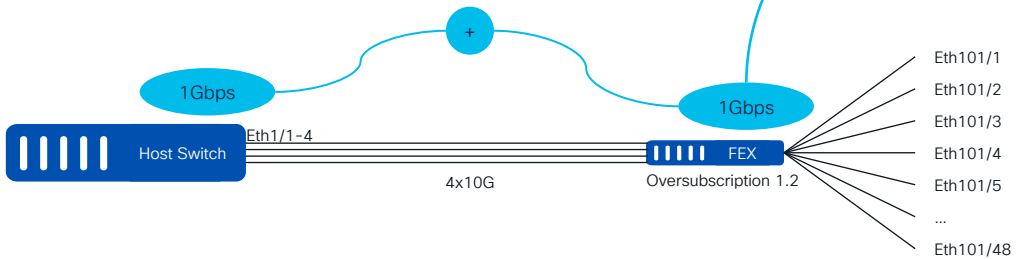

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Pricing Economics

Nexus 9000 + FEX

Around 2013

Host Switch	Price USD$	$ per Gbps
Cisco Nexus 93180YC-FX • 48 x 10G + 6 x 100G	USD $30,400.00	USD ~$28.00/Gbps
(48*10Gbps)+(6*100Gbps) =1080Gbps/Switch	USD $30,400.00/1080Gbps	

FEX	Price USD$	$ per Gbps
Cisco Nexus 2348UPQF-QSA • 48 x 1G + 6 x 10G	USD $24,300.00	USD $506.00/Gbps
48*1Gbps =48Gbps HIF per FEX	USD $24,300.00/48Gbps	

FEX Port $ per Gbps (1.2 Oversubscription)

USD ~$534.00/Gbps

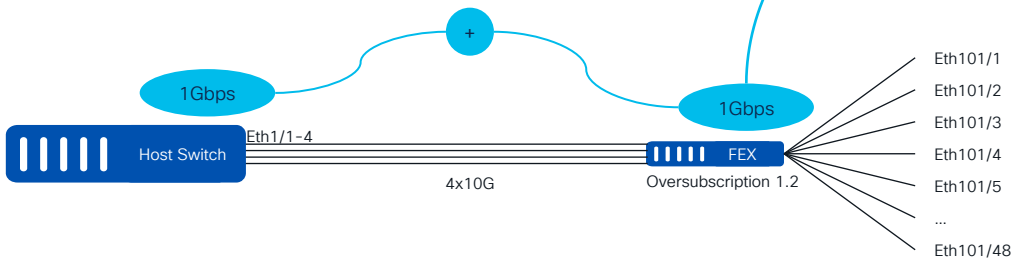

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Pricing Economics

Nexus 9000 + ToR

Today

Host Switch	Price USD$	$ per Gbps
Cisco Nexus 93180YC-FX • 48 x 10G + 6 x 100G	USD $30,400.00	USD ~$28.00/Gbps
(48*10Gbps)+(6*100Gbps) =1080Gbps/Switch	USD $30,400.00/1080Gbps	

TOR	Price USD$	$ per Gbps
Cisco Nexus 9348GC-FXP • 48 x 1G + 2 x 100G	USD $17,500.00	USD ~$365.00/Gbps
48*1Gbps =48Gbps HIF per TOR	USD $17,500.00/48Gbps	

TOR Port $ per Gbps (1.2 Oversubscription)

USD ~$393.00/Gbps

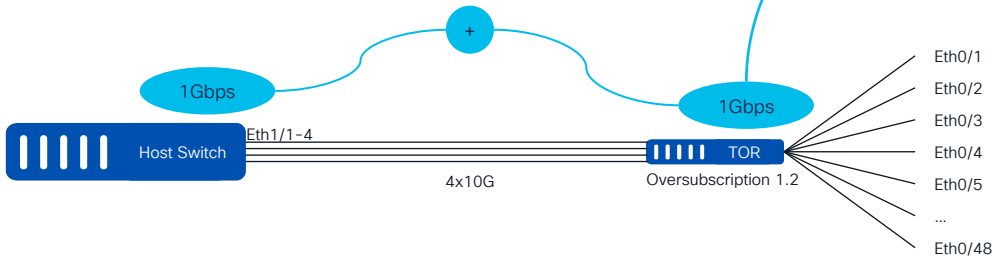

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Pricing Economics Comparison

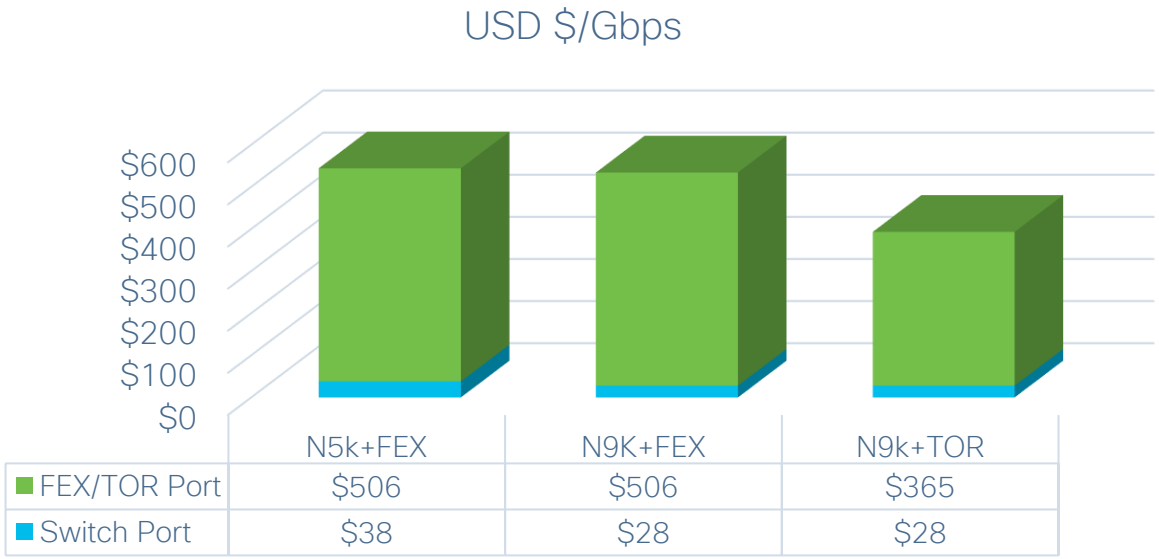

Bandwidth/Cost Change over a Decade

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Pricing Economics

Nexus 9000 + FEX with 1536 Ports @1Gbps

Today

Host Switch	Price USD$ * 4 Host Switch	$ per Gbps
Cisco Nexus 93180YC-FX • 48 x 10G + 6 x 100G	USD $30,400.00	USD ~$28.00/Gbps
(48*10Gbps)+(6*100Gbps) =1080Gbps/Switch	USD $30,400.00/1080Gbps	

FEX	Price USD$ * 32 FEX	$ per Gbps
Cisco Nexus 2348UPQF-QSA • 48 x 1G + 6 x 10G	USD $24,300.00	USD $506.00/Gbps
48*1Gbps =48Gbps HIF per FEX	USD $24,300.00/48Gbps	

FEX Port $ per Gbps

USD ~$534.00/Gbps

USD ~$585.00/Host Port

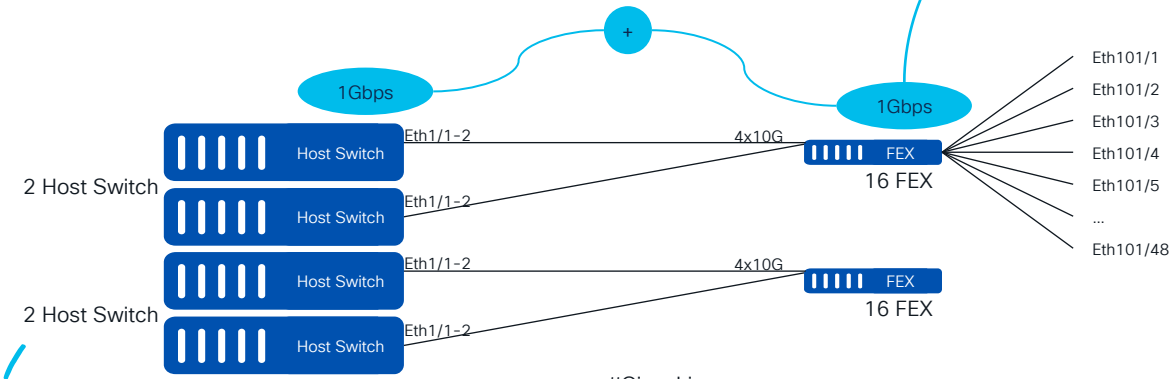

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Pricing Economics

Nexus 9000 Fabric with 1536 Ports @1Gbps

Today

Spine	Price USD$ * 2 Spine	$ per Gbps
Cisco Nexus 9336C-FX2 • 36 x 100G	USD $38,800.00	USD ~$11.00/Gbps
36*100Gbps =3'600Gbps/Switch	USD $38,800.00/3600Gbps	

Leaf	Price USD$ * 32 Leaf	$ per Gbps
Cisco Nexus 9348GC-FXP • 48 x 1G + 2x 100G	USD $17,500.00	USD $365.00/Gbps
48*1Gbps =48Gbps HIF per TOR	USD $12,000.00/48Gbps	

TOR Port $ per Gbps	
USD ~$376.00/Gbps	USD ~$415.00/Host Port

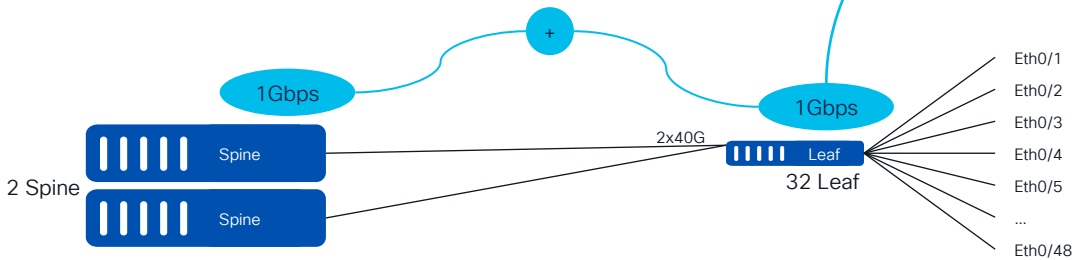

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Pricing Economics Comparison with 1536 Ports @1Gbps

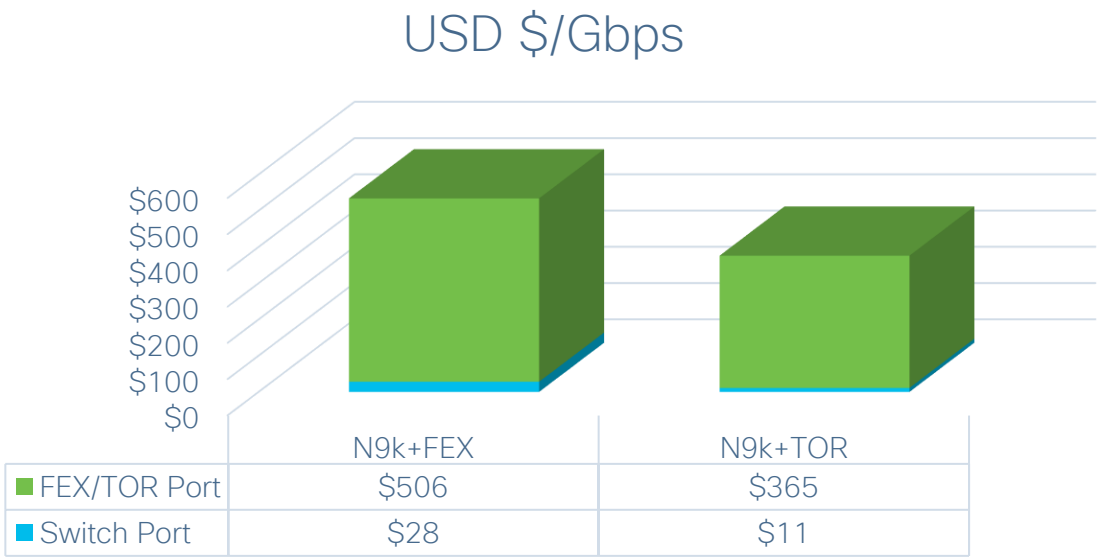

	N9k+FEX	N9k+TOR
FEX/TOR Port	$506	$365
Switch Port	$28	$11

Optimizing Further with Port Count

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration Considerations

Migration Considerations

The Usual Approach of Building a New Parallel Network (1)

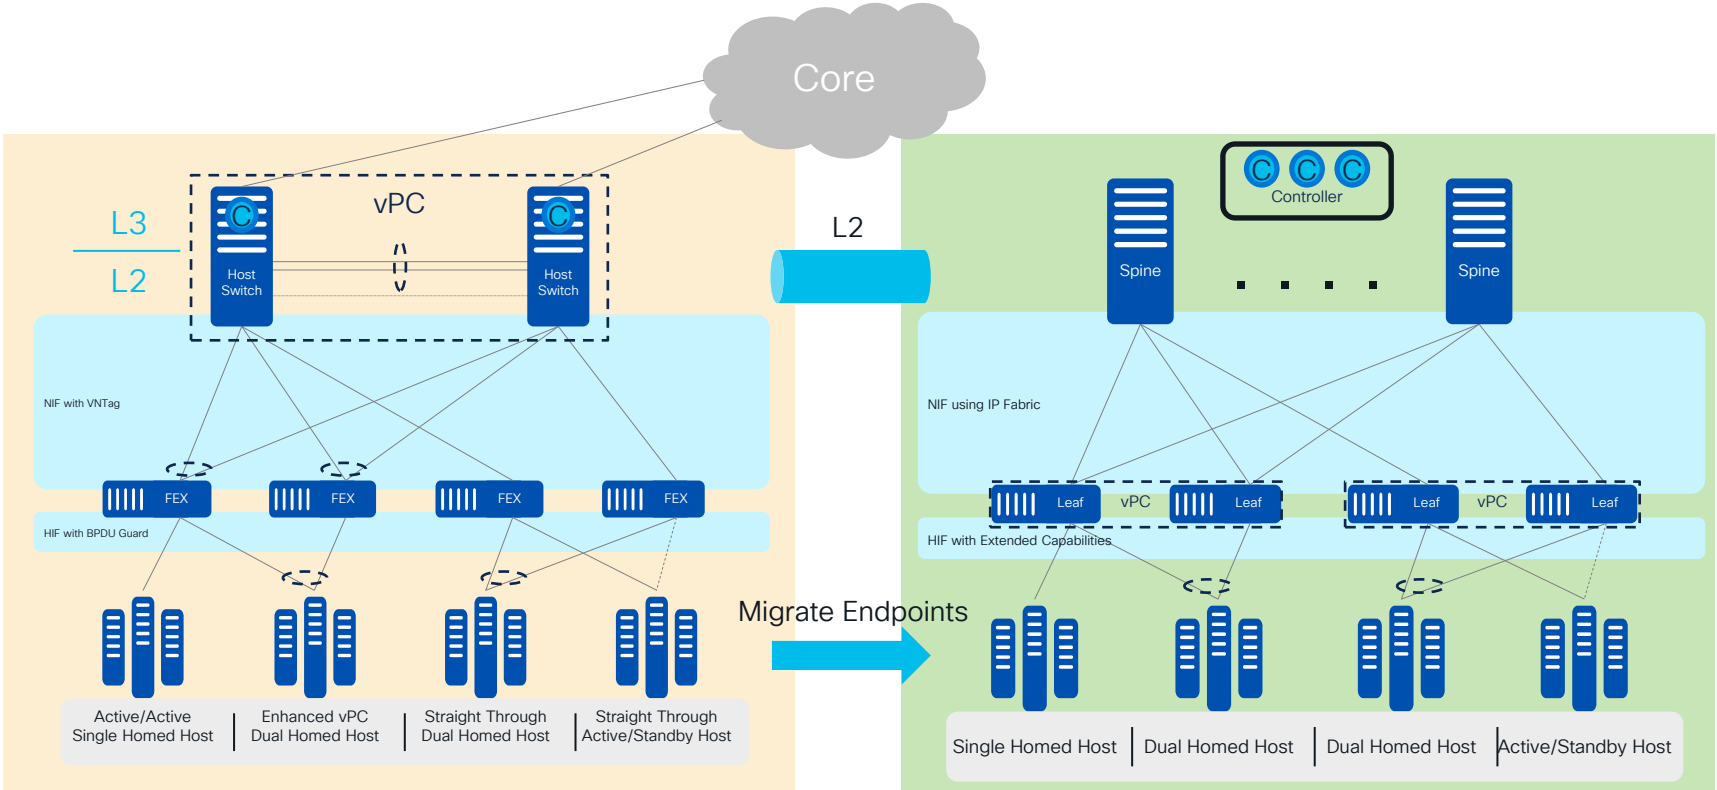

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration Considerations

The Usual Approach of Building a New Parallel Network (2)

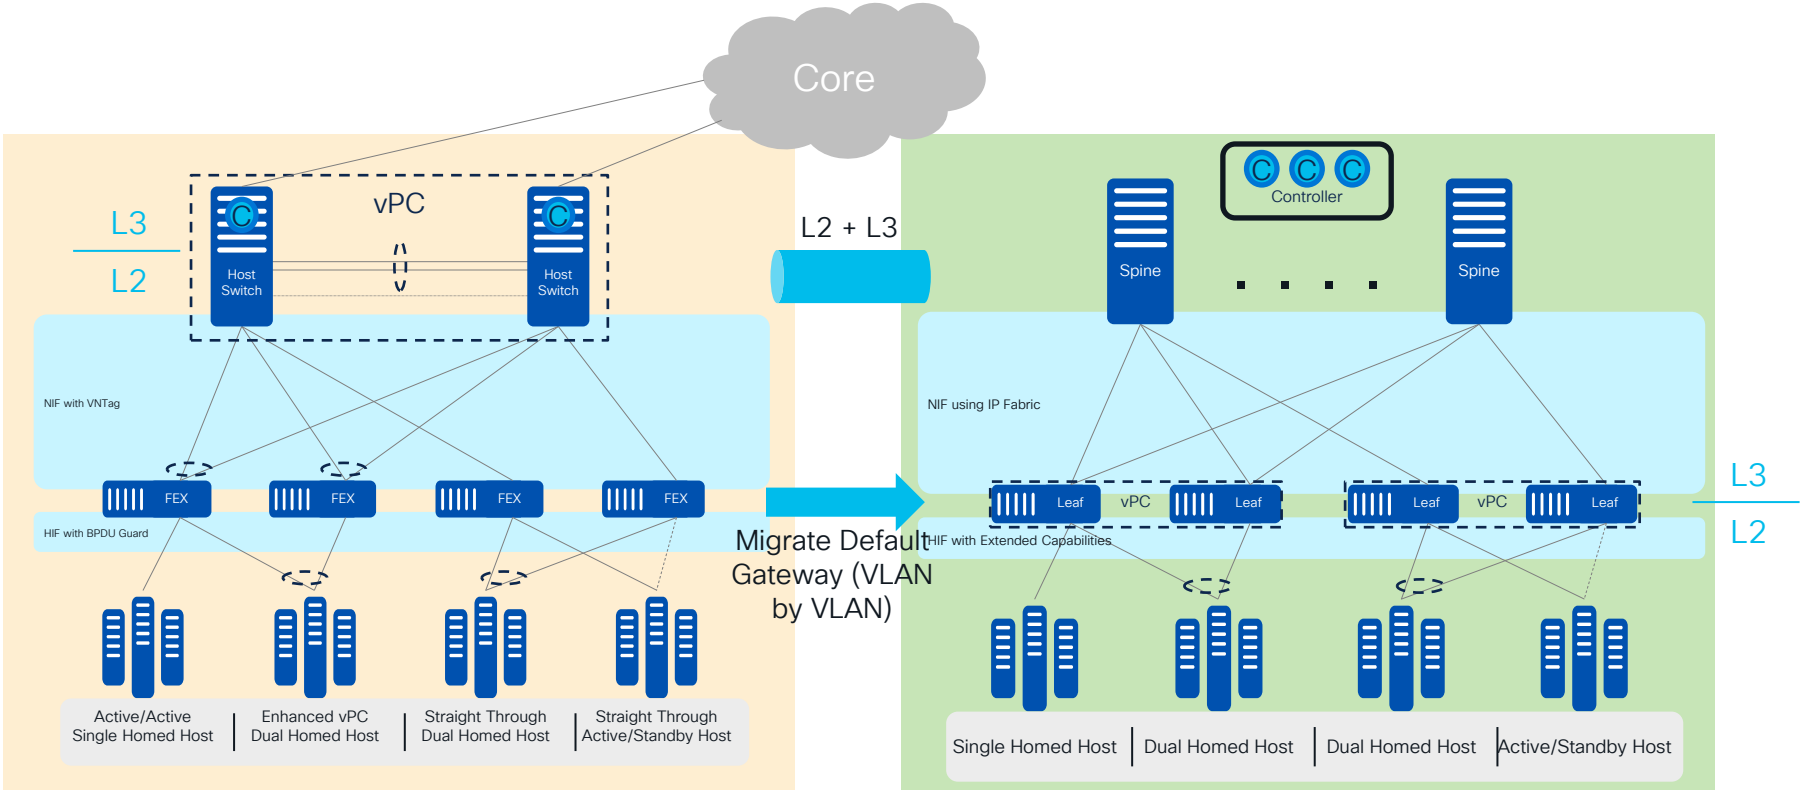

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration Considerations

The Usual Approach of Building a New Parallel Network (3)

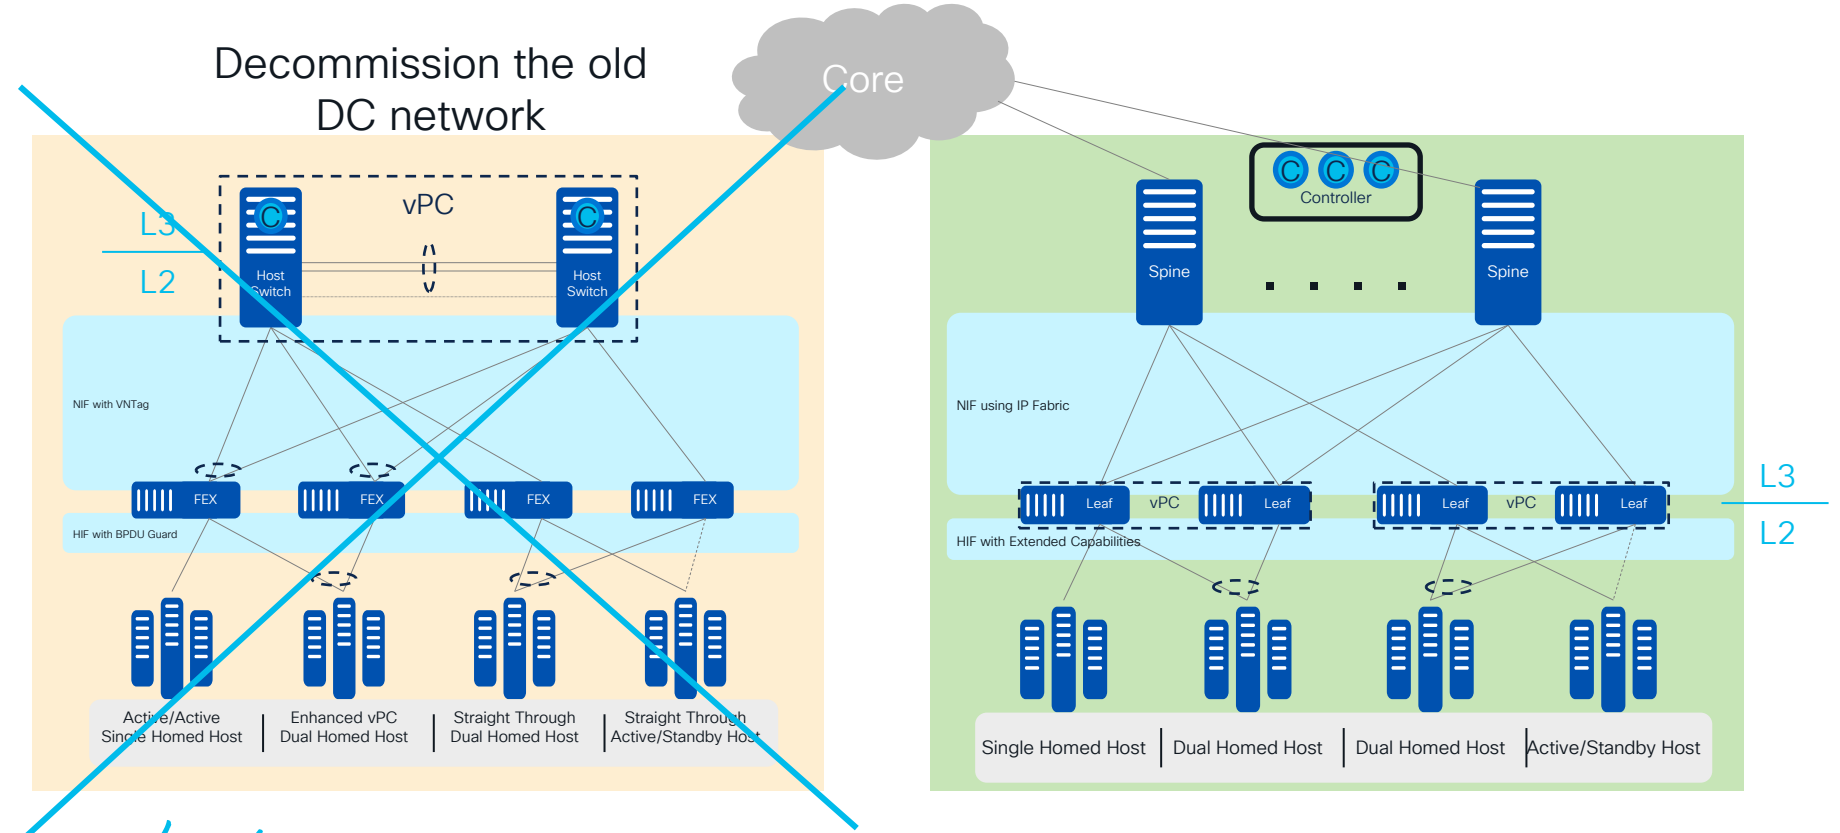

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration with Rack Space Constraints

Migration with Rack Space Constraints

Starting Point

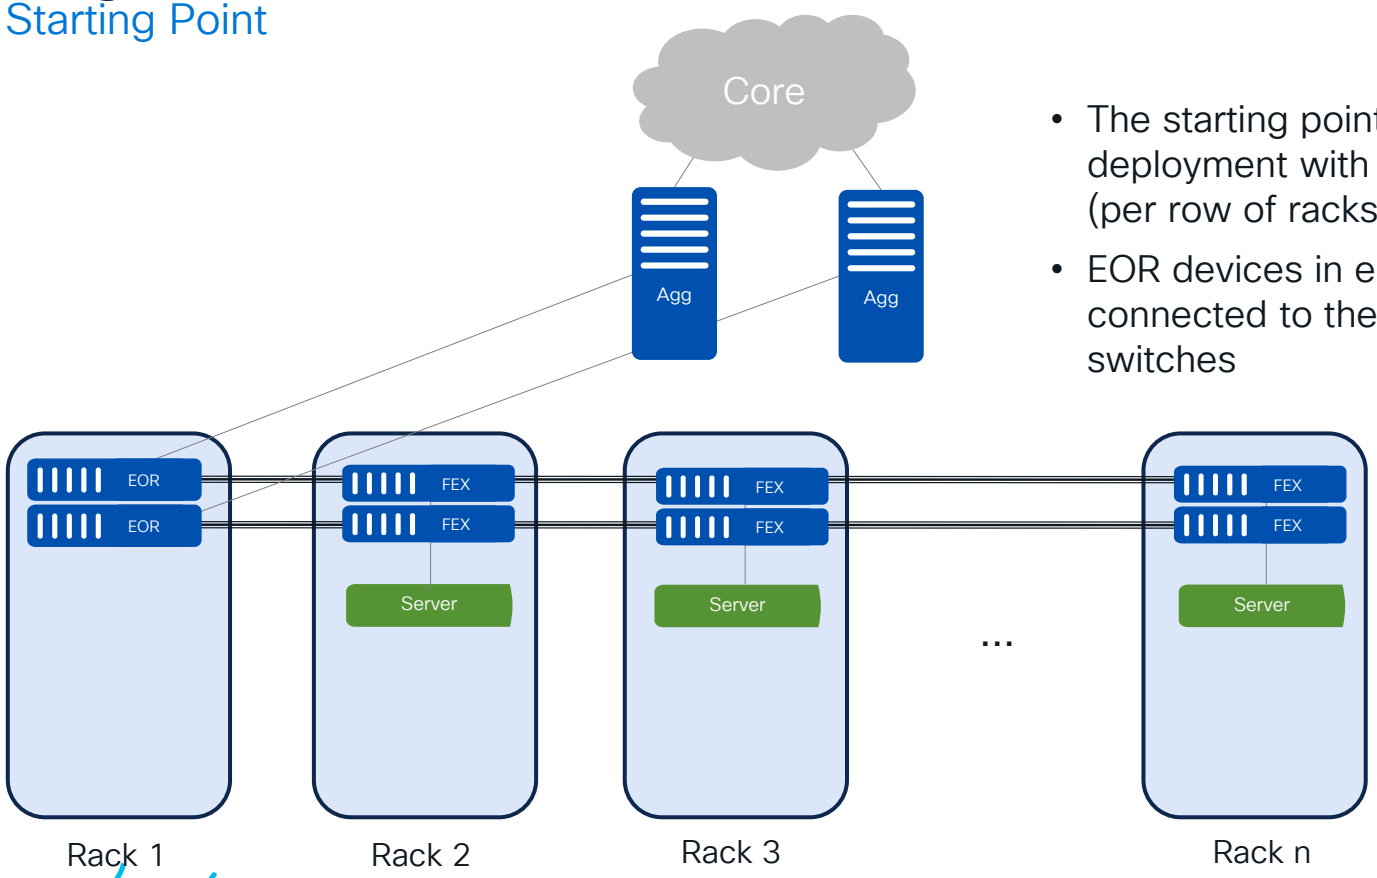

- The starting point is the traditional FEX deployment with a pair of EOR devices (per row of racks)
- EOR devices in each row are connected to the centralized Agg switches

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration with Rack Space Constraints

Adding New Spine/Leaf/BL Nodes

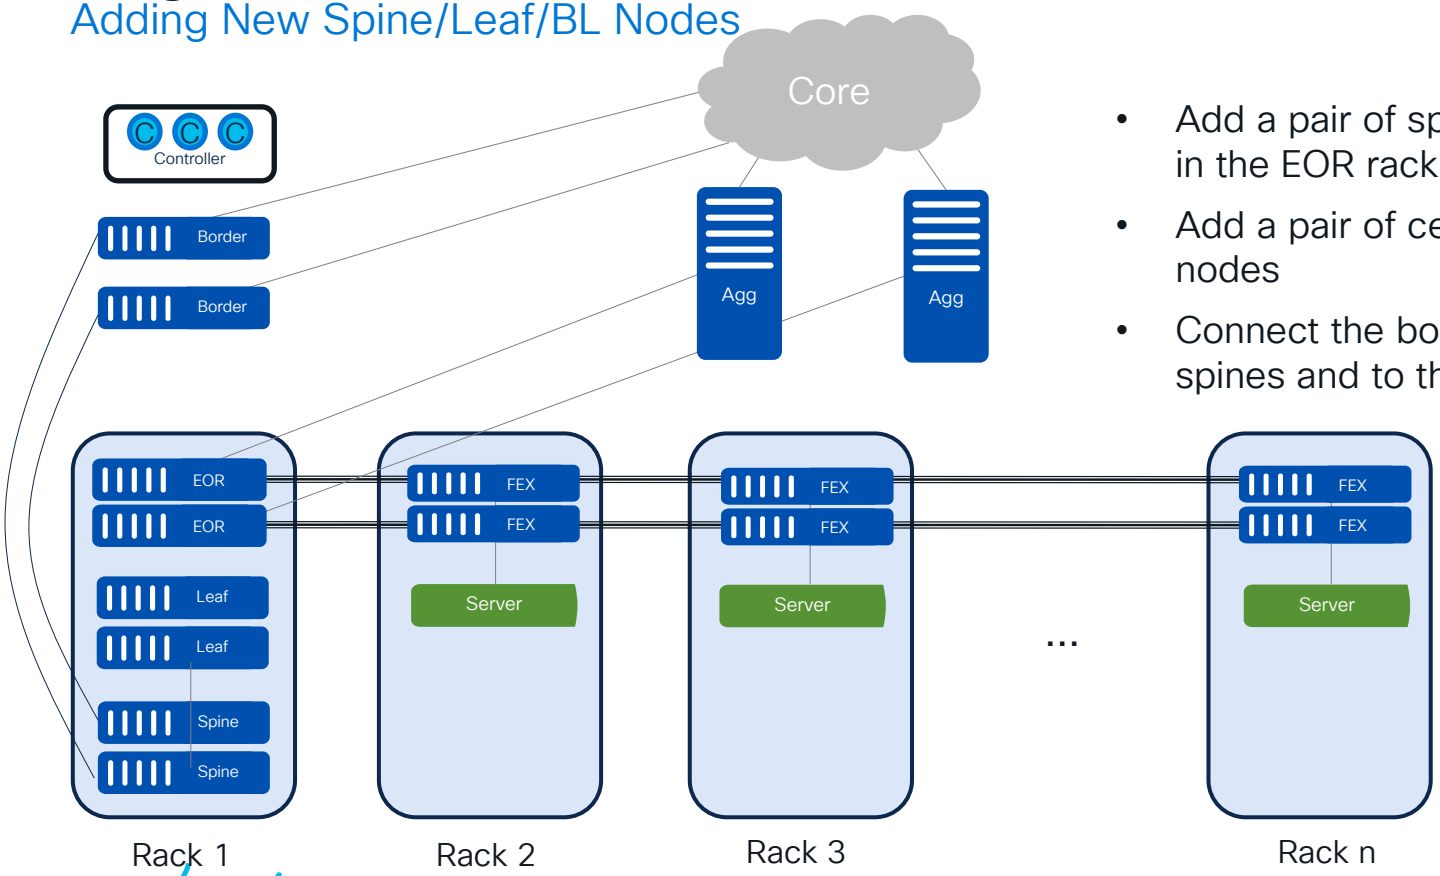

- Add a pair of spine and leaf switches in the EOR rack
- Add a pair of centralized border nodes
- Connect the border nodes to the spines and to the core

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration with Rack Space Constraints

Adding New Spine/Leaf/BL Nodes

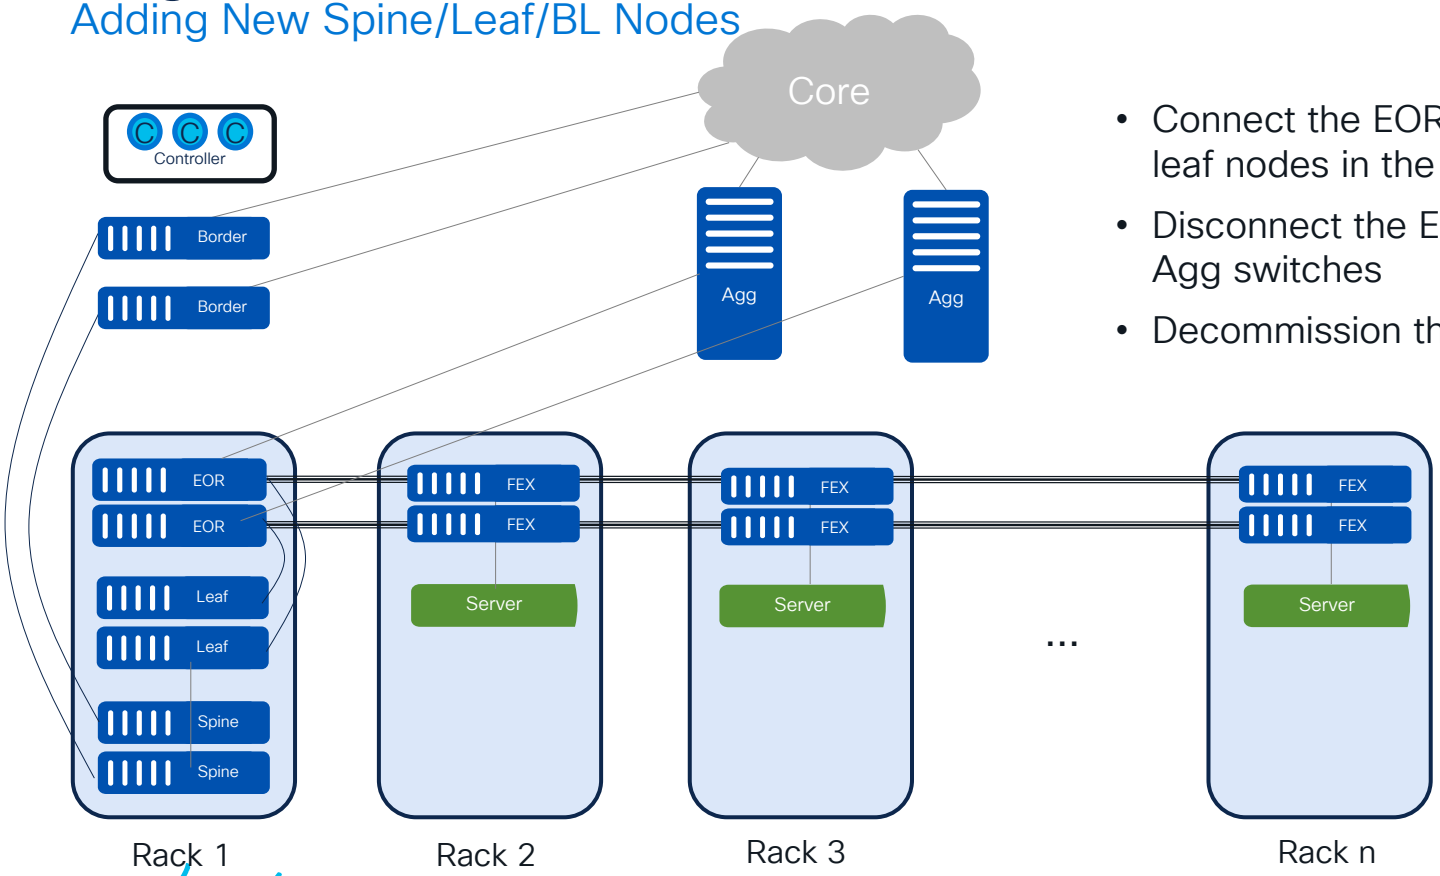

- Connect the EOR devices to the pair of leaf nodes in the EOR rack (L2 + L3)
- Disconnect the EOR devices from the Agg switches
- Decommission the Agg switches

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration with Rack Space Constraints

Adding New Spine/Leaf/BL Nodes

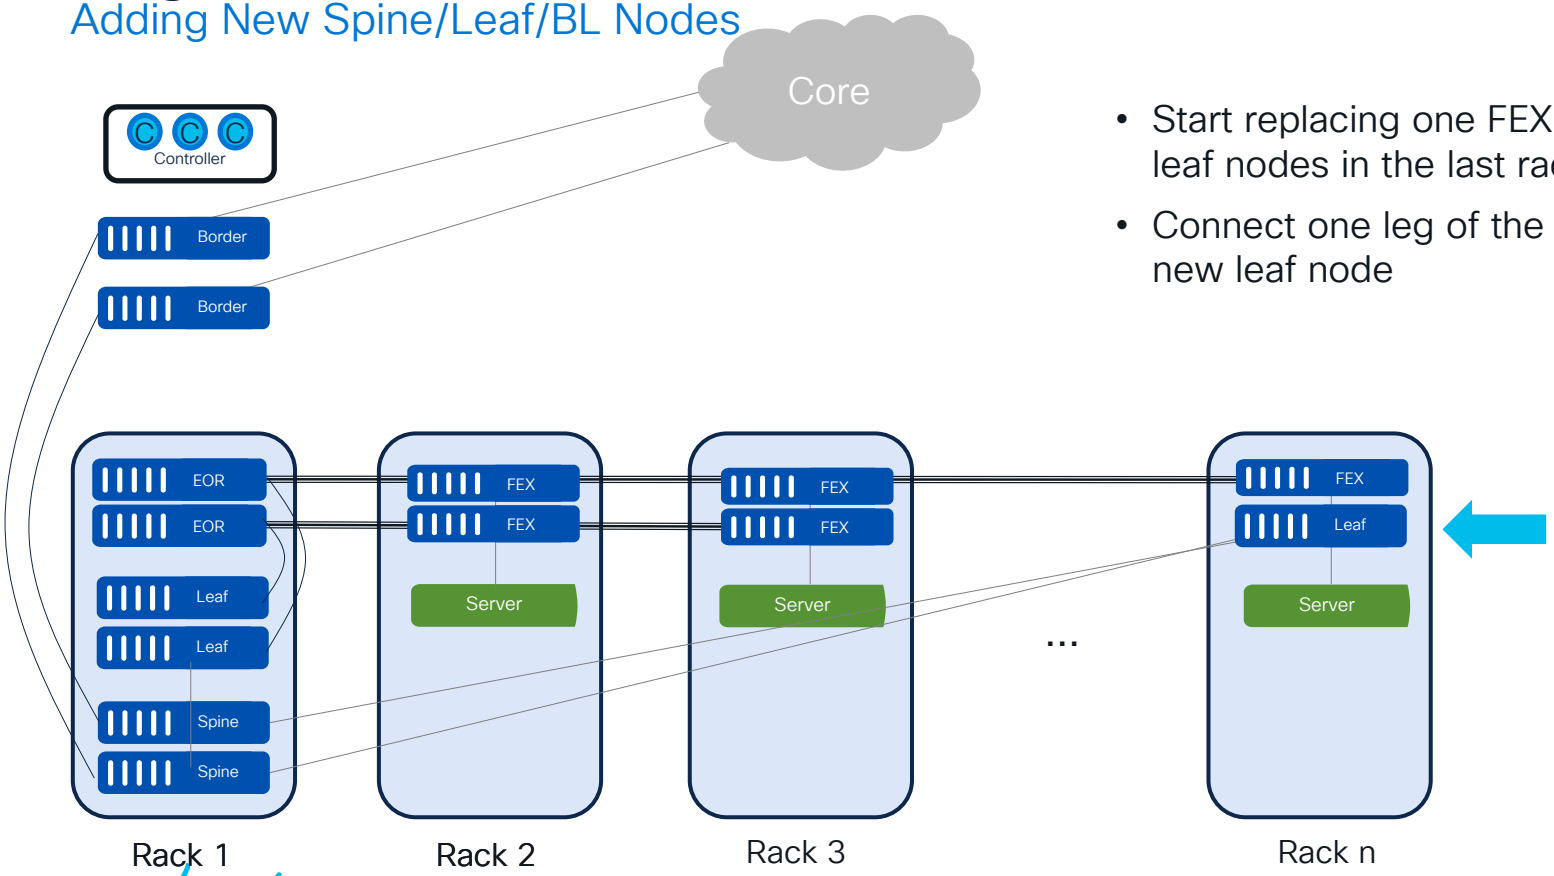

- Start replacing one FEX with a new leaf nodes in the last rack
- Connect one leg of the servers to the new leaf node

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration with Rack Space Constraints

Adding New Spine/Leaf/BL Nodes

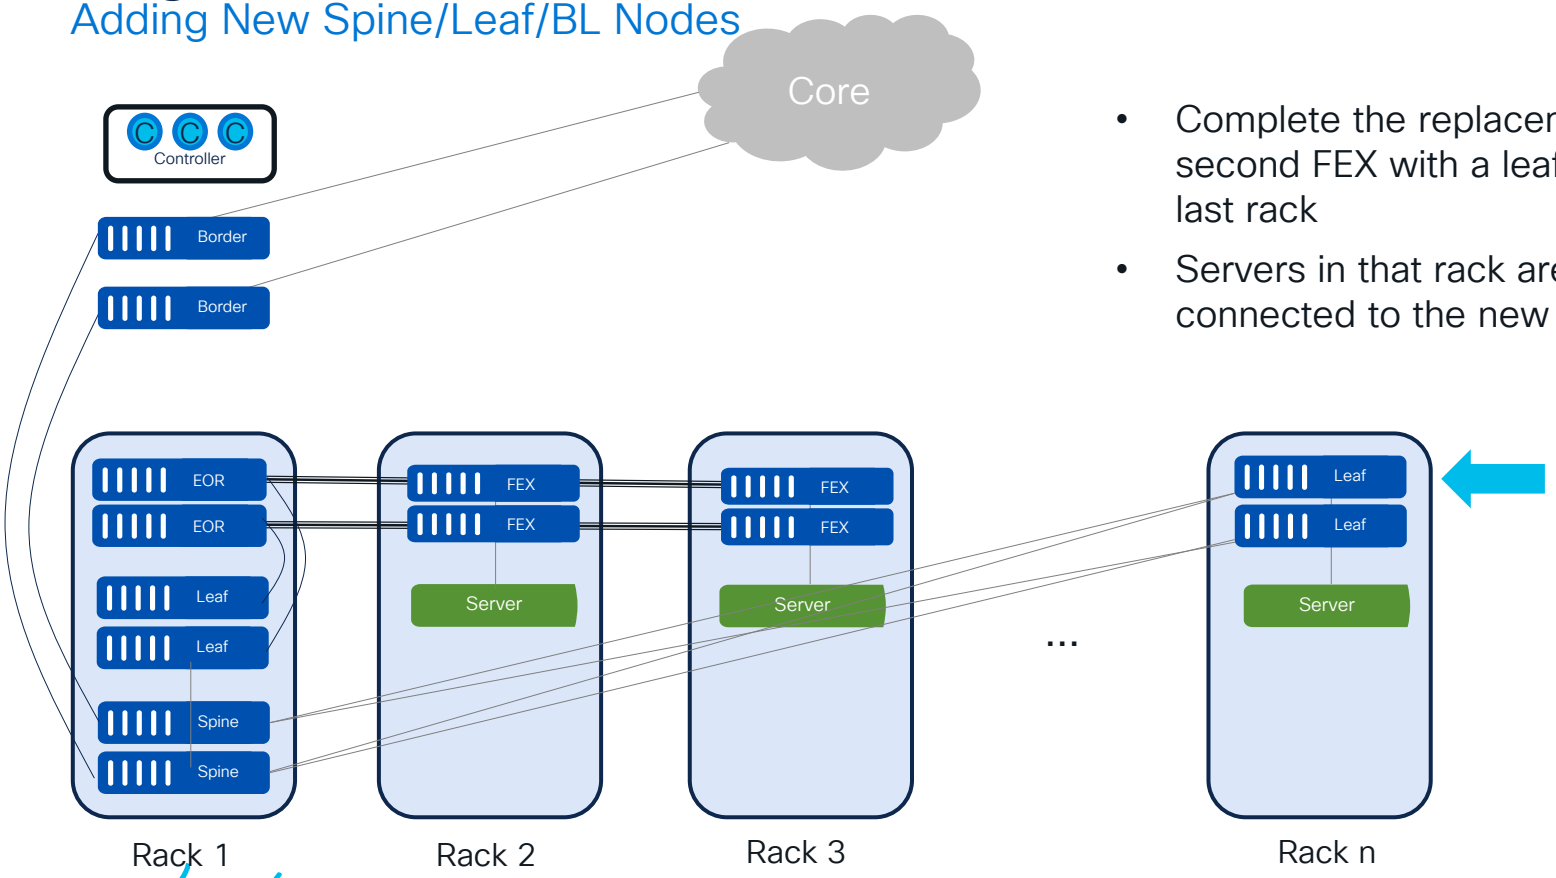

- Complete the replacement of the second FEX with a leaf node in the last rack
- Servers in that rack are now only connected to the new fabric

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration with Rack Space Constraints

Adding New Spine/Leaf/BL Nodes

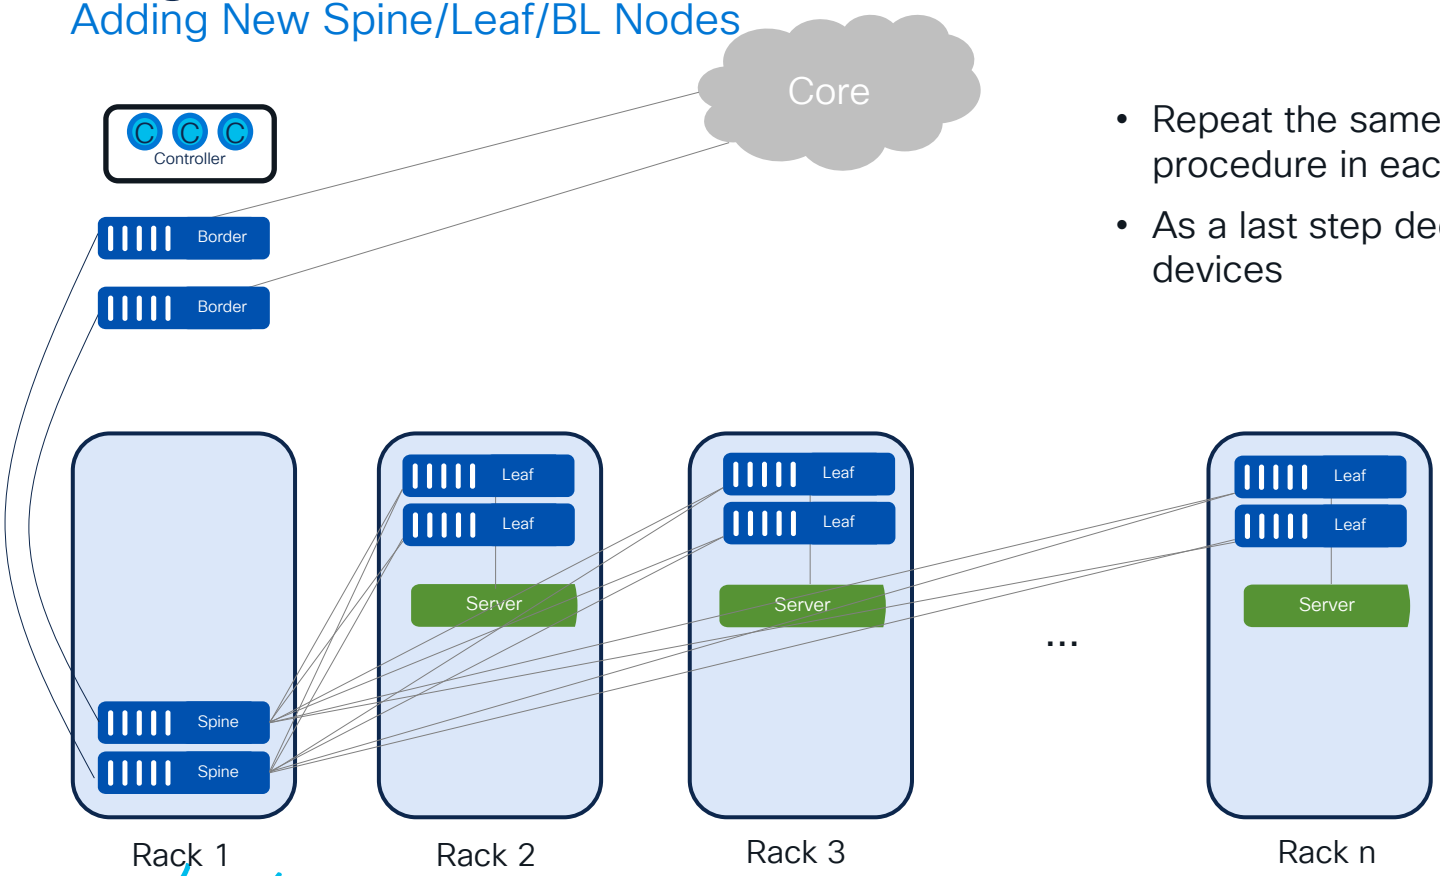

- Repeat the same FEX replacement procedure in each rack
- As a last step decommission the EOR devices

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration without Rack Space Constraints

Migration without Rack Space Constraints

Starting Point

1

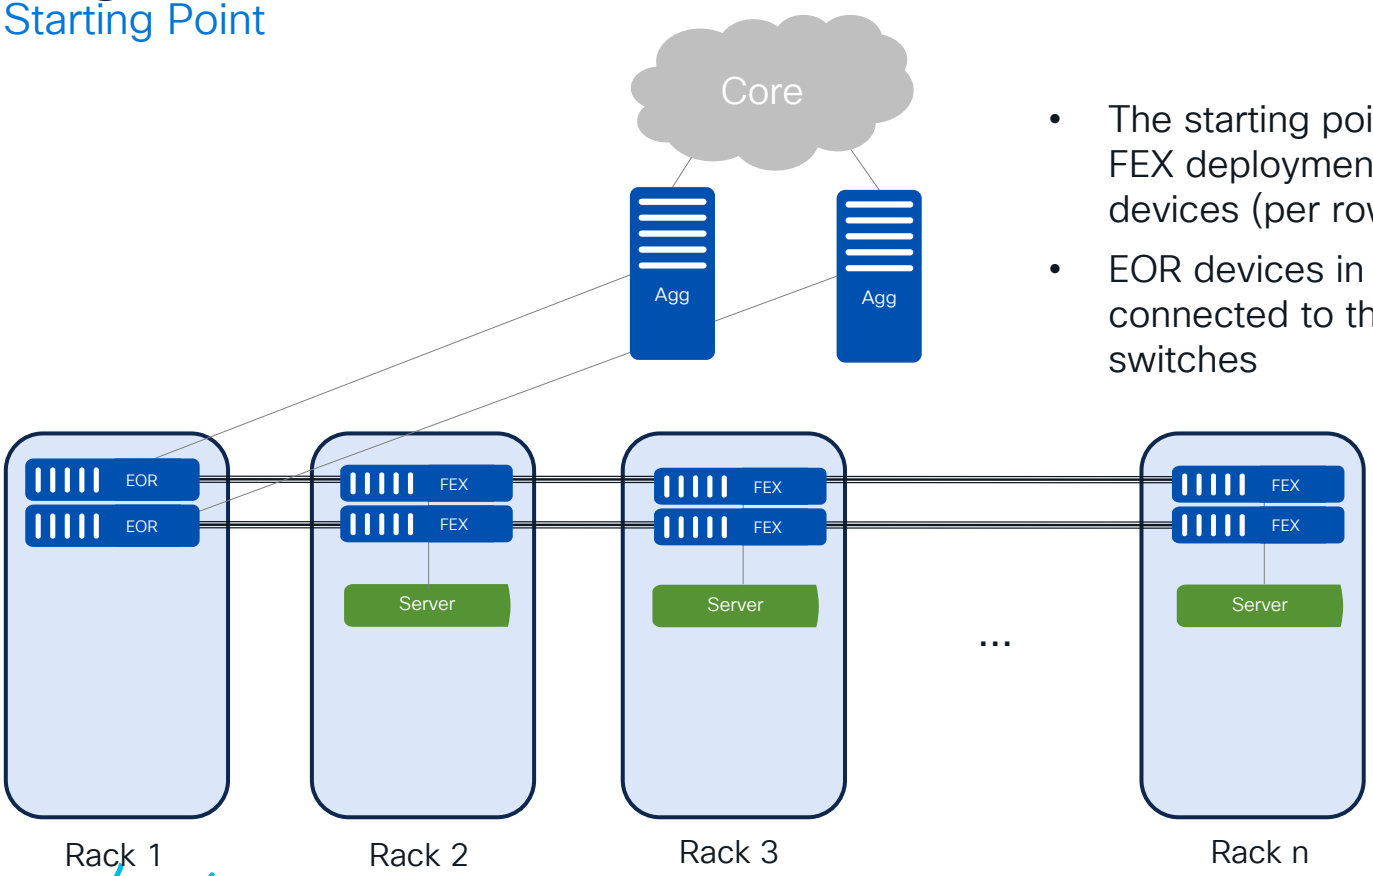

- The starting point is the traditional FEX deployment with a pair of EOR devices (per row of racks)
- EOR devices in each row are connected to the centralized Agg switches

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration without Rack Space Constraints

Adding New Spine/Leaf/BL Nodes

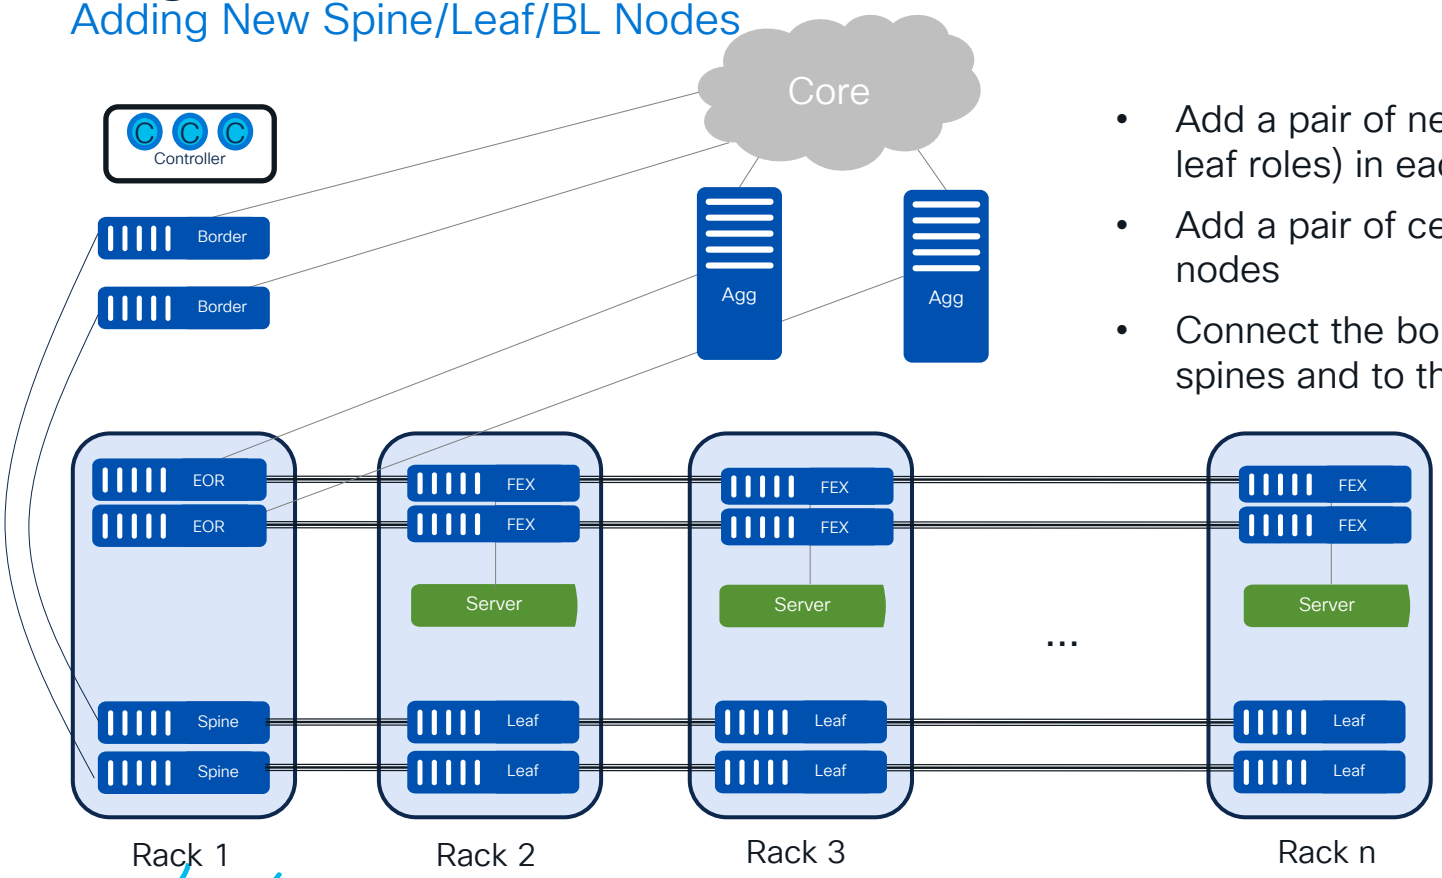

- Add a pair of new devices (spine and leaf roles) in each rack
- Add a pair of centralized border nodes
- Connect the border nodes to the spines and to the core

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration without Rack Space Constraints

Connect Old and New Infrastructures and Migrate Endpoints

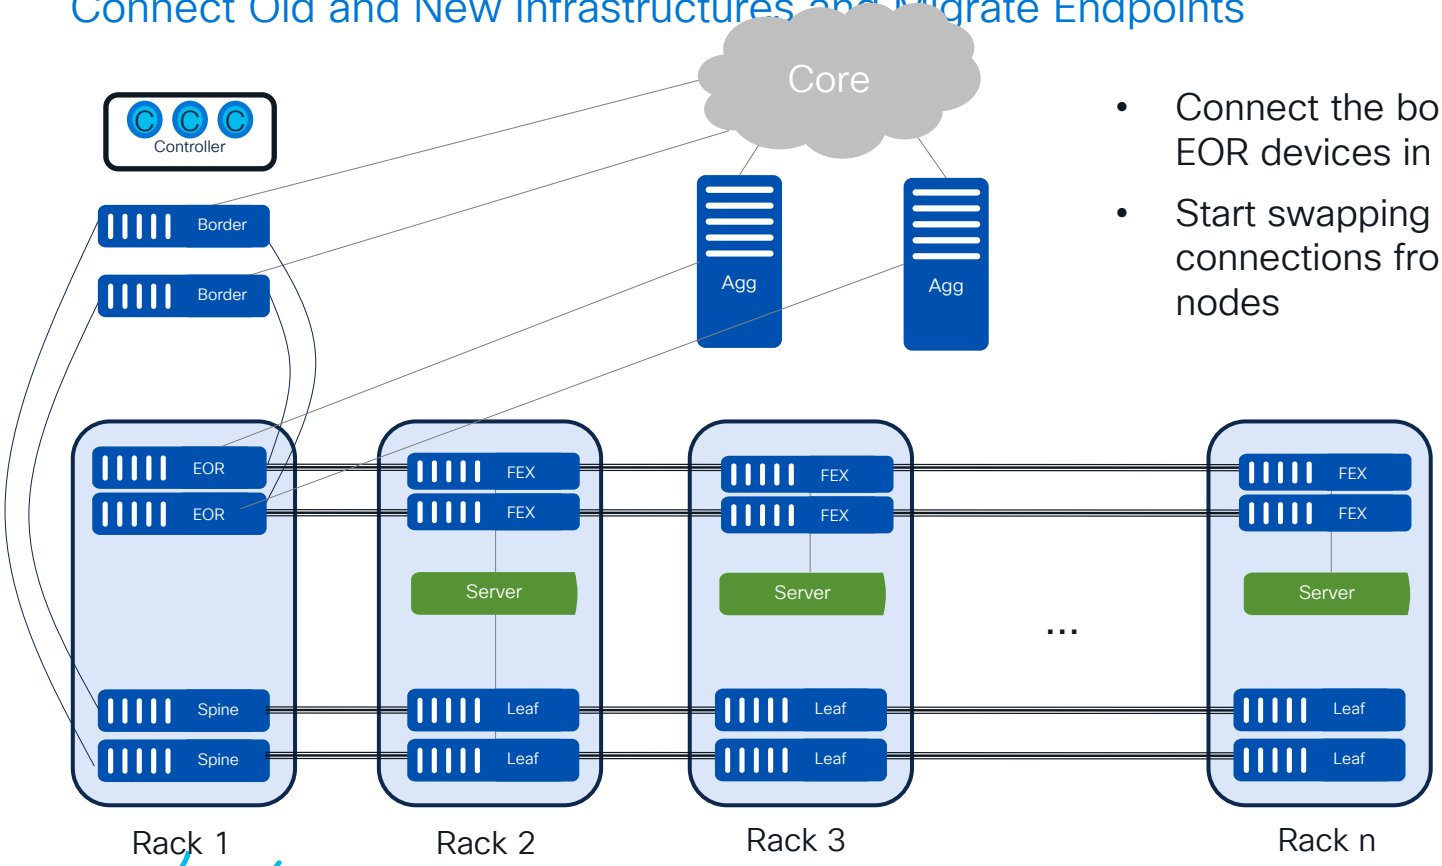

- Connect the border nodes to the EOR devices in each row (L2 and L3)
- Start swapping the servers connections from the FEX to the leaf nodes

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Migration without Rack Space Constraints

Decommission the Old Infrastructure

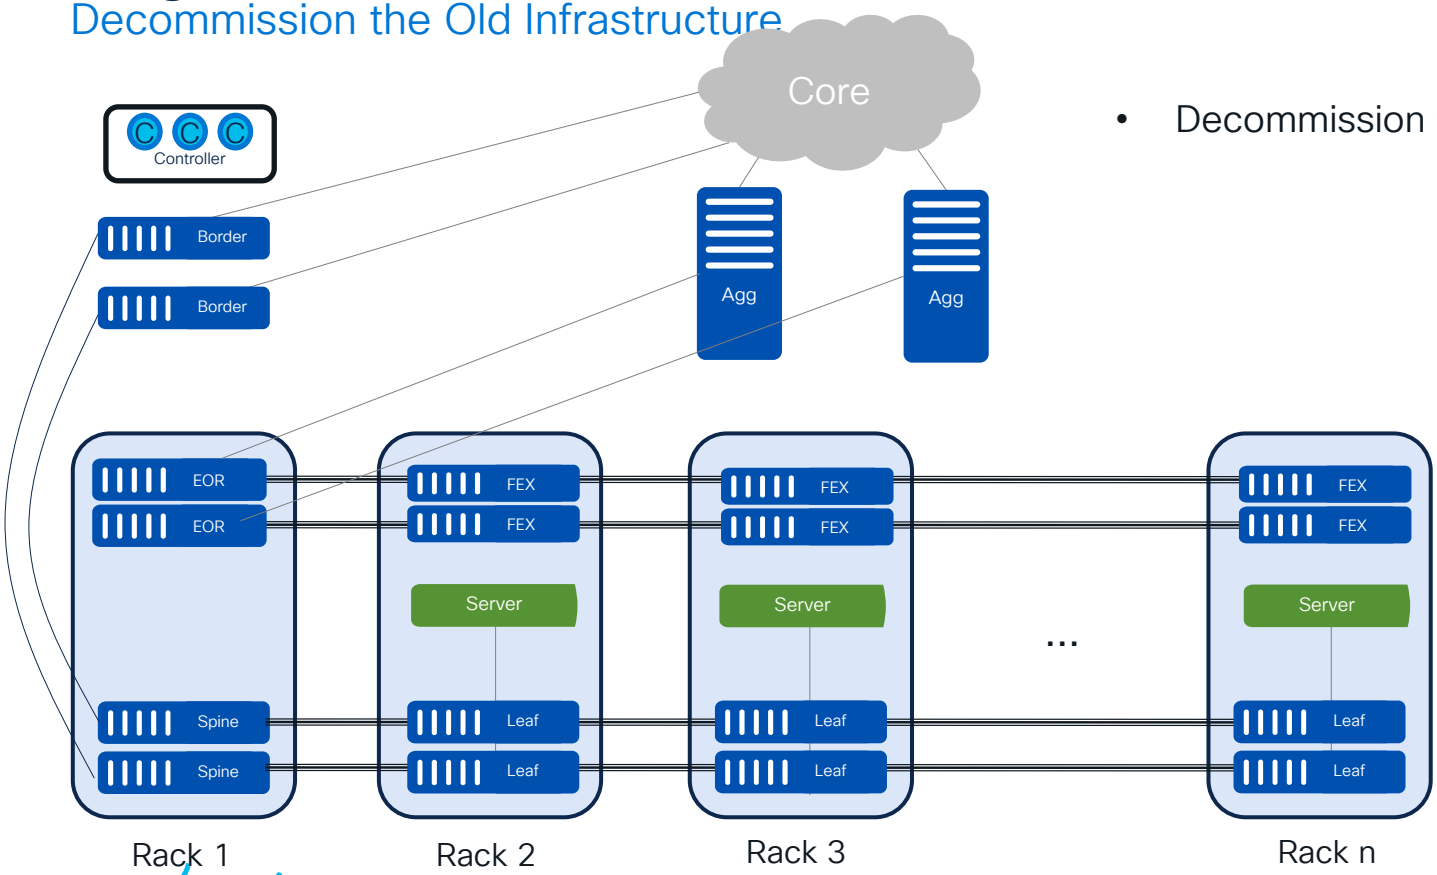

- Decommission the old infrastructure

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Conclusion

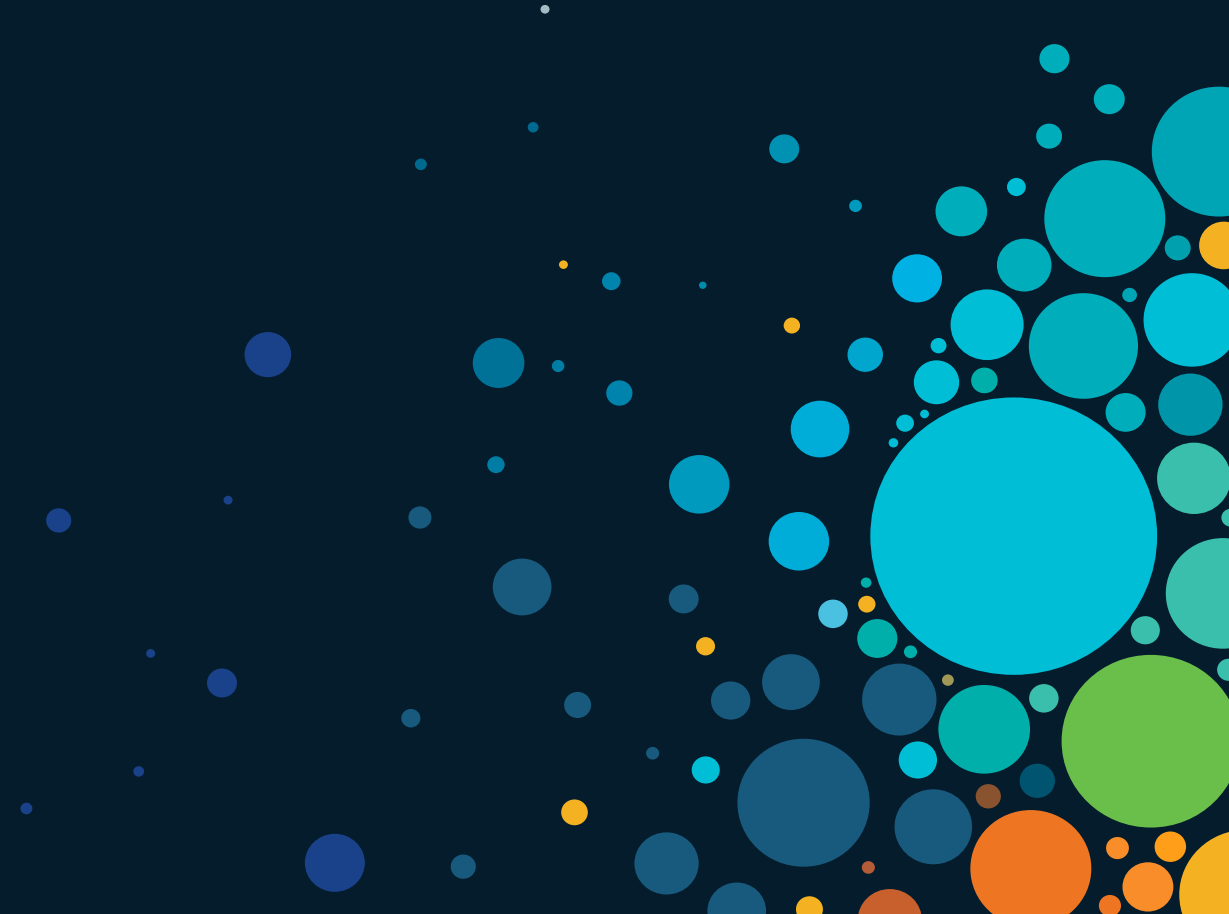

Conclusions

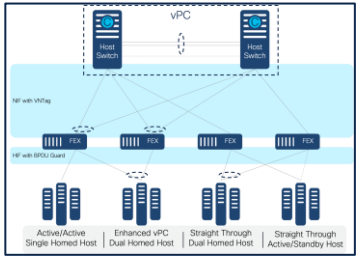

- FEX was the first attempt to build a fabric infrastructure
  - Centralized Management
  - Network and Host Redundancy

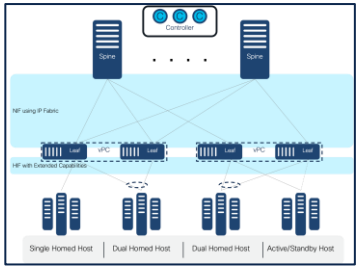

- Evolution of network architectures to deliver full fledge fabrics
  - Centralized Management with Controller
  - Fully distributed control and data planes

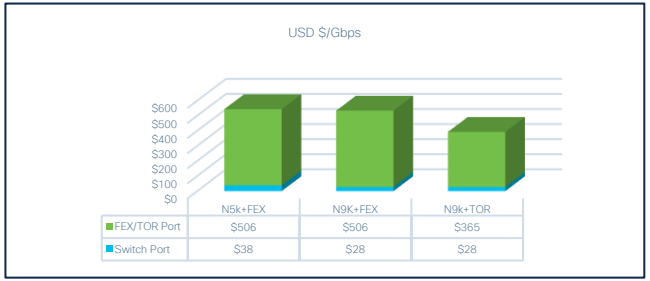

- Bandwidth/Cost Evolution over a Decade
- Economics started favoring deployment of switches as ToRs

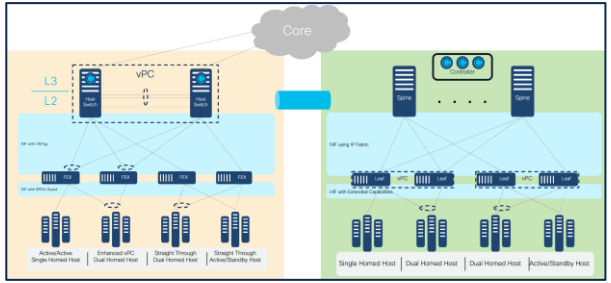

- Usual migration approach of building a parallel network
- Couple options based on existence of rack space constraints

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

Learn

Cisco U.
IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning
Subscription-based product, technology, and certification training

Cisco Modeling Labs
Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network
Resource community portal for certifications and learning

Train

Cisco Training Bootcamps
Intensive team & individual automation and technology training programs

Cisco Learning Partner Program
Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training
Accelerated curriculum of product, technology, and certification courses

Certify

Cisco Certifications and Specialist Certifications
Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups
180-day certification prep program with learning and support

Cisco Continuing Education Program
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

cisco Live!

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Public

Thank you


cisco Live!
ALL IN