



**OPEN**  
Compute  
Project®

# Smart Networking Needed at Every Tier

David Iles

Sr. Director at NVIDIA

Connect. Collaborate. Accelerate.

# The Problem

- Compromises were made:

SPIN


Optimized for Speed

Dumb/Fast Switches

LEAF

Optimized for  
Features

Smart/Slow Switches



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# The Problem

- Compromises were made:

SPIN

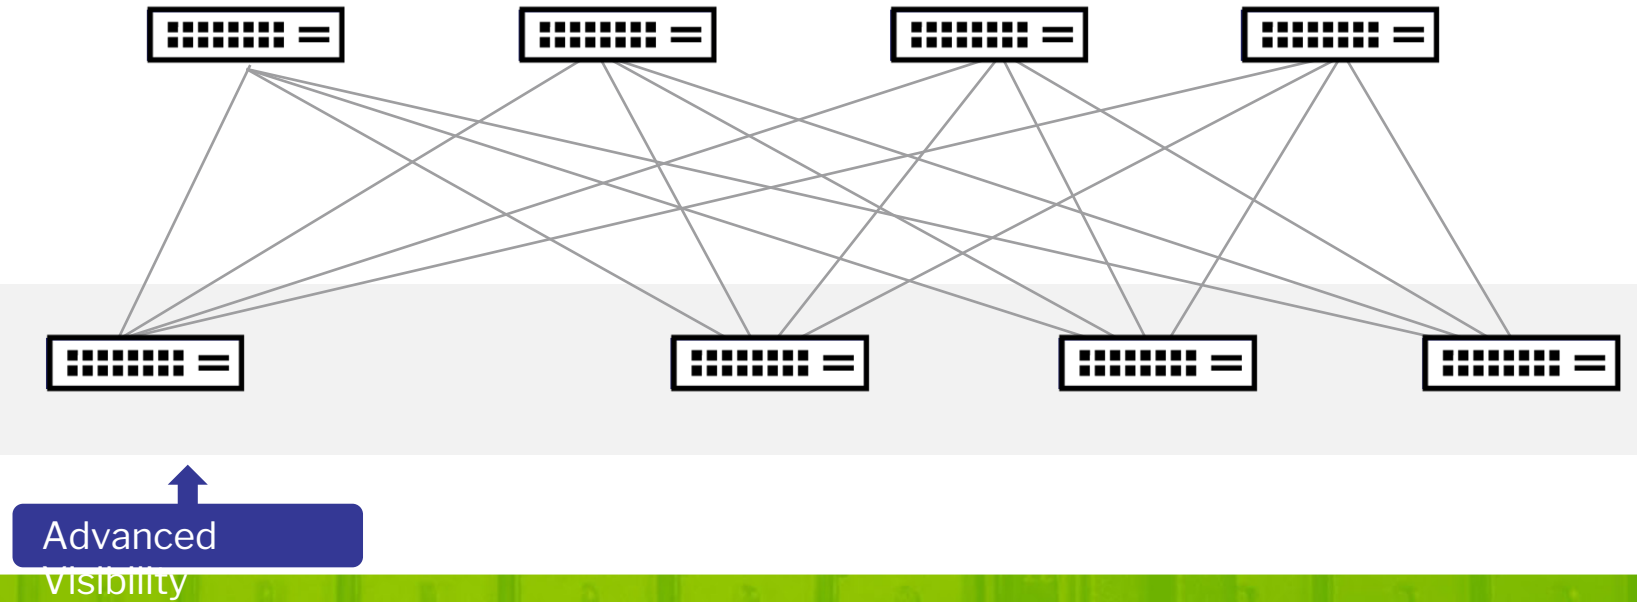
Optimized for Speed

Dumb/Fast Switches

LEAF

Optimized for Features

Smart/Slow Switches



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# The Problem

- Compromises were made:

SPIN

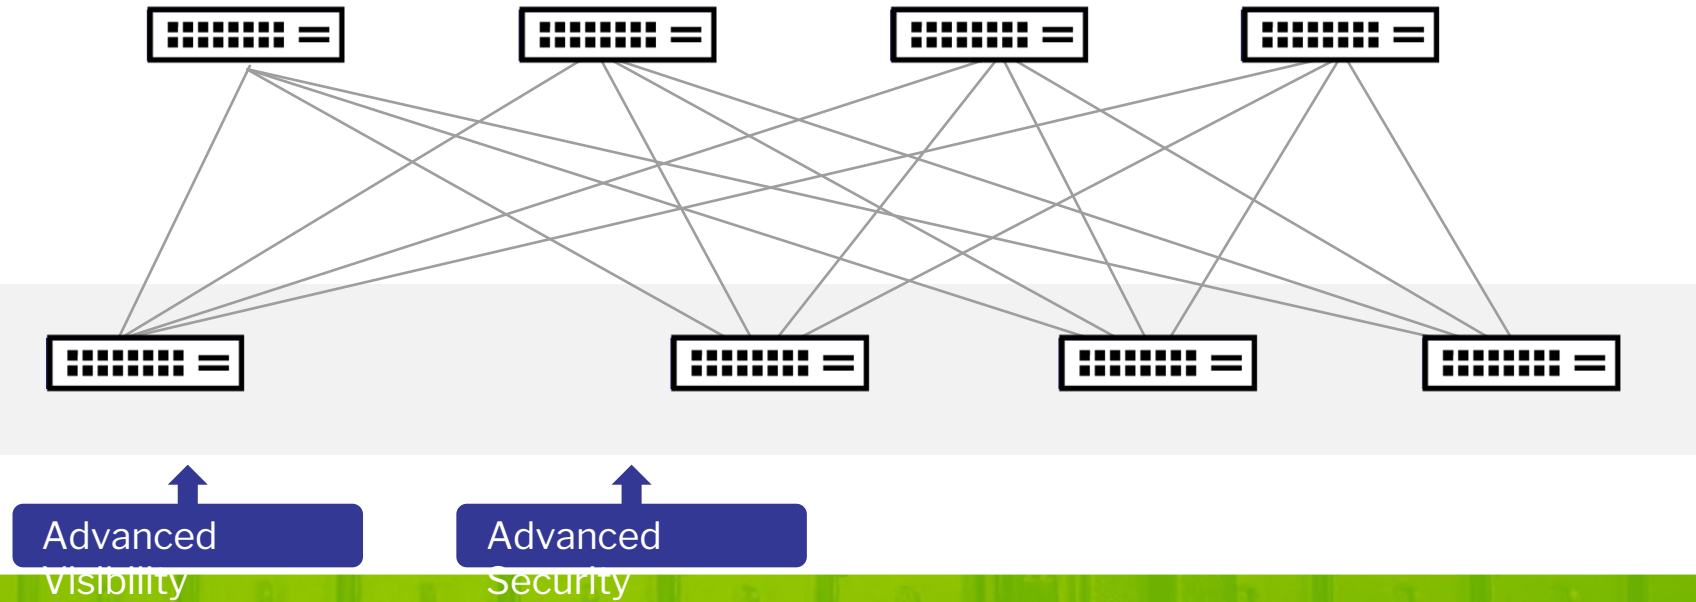
Optimized for Speed

Dumb/Fast Switches

LEAF

Optimized for Features

Smart/Slow Switches



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.



# The Problem

- Compromises were made:

SPIN

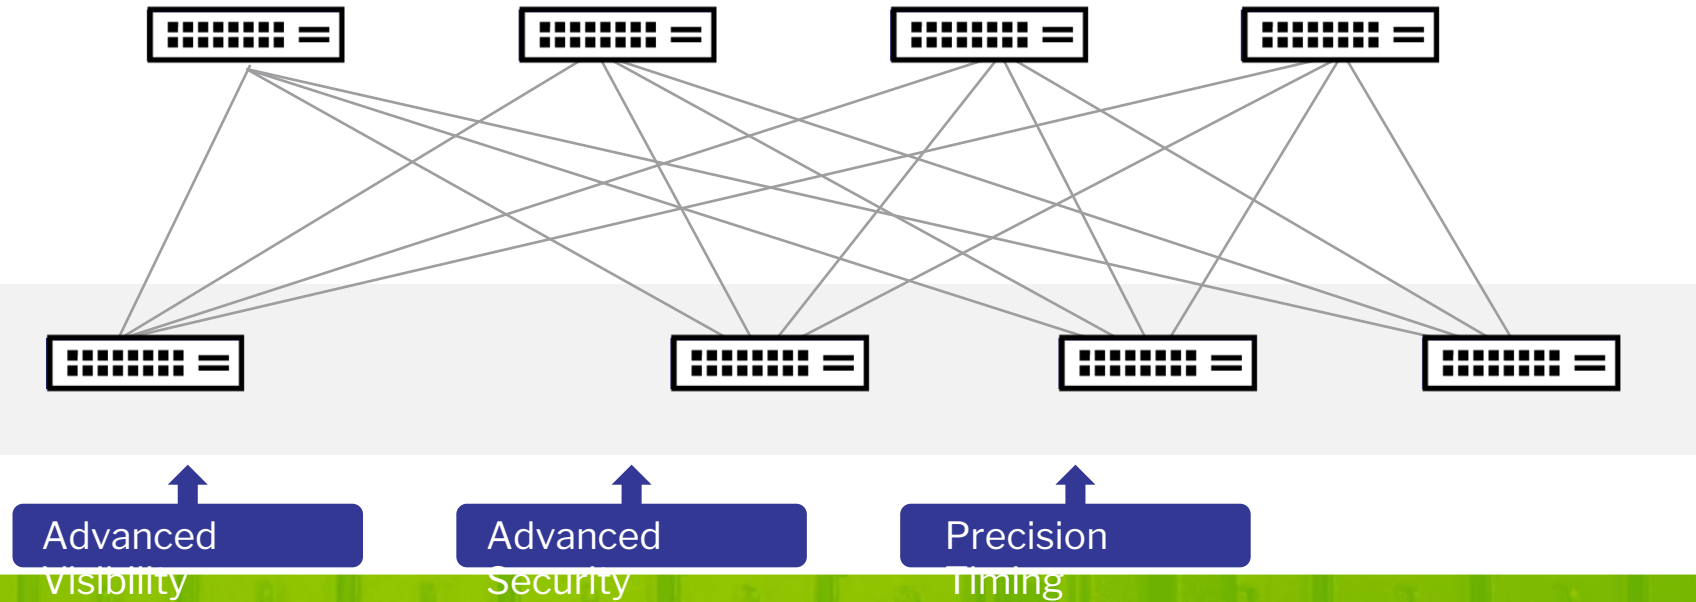
Optimized for Speed

Dumb/Fast Switches

LEAF

Optimized for Features

Smart/Slow Switches



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# The Problem

- Compromises were made:

SPIN

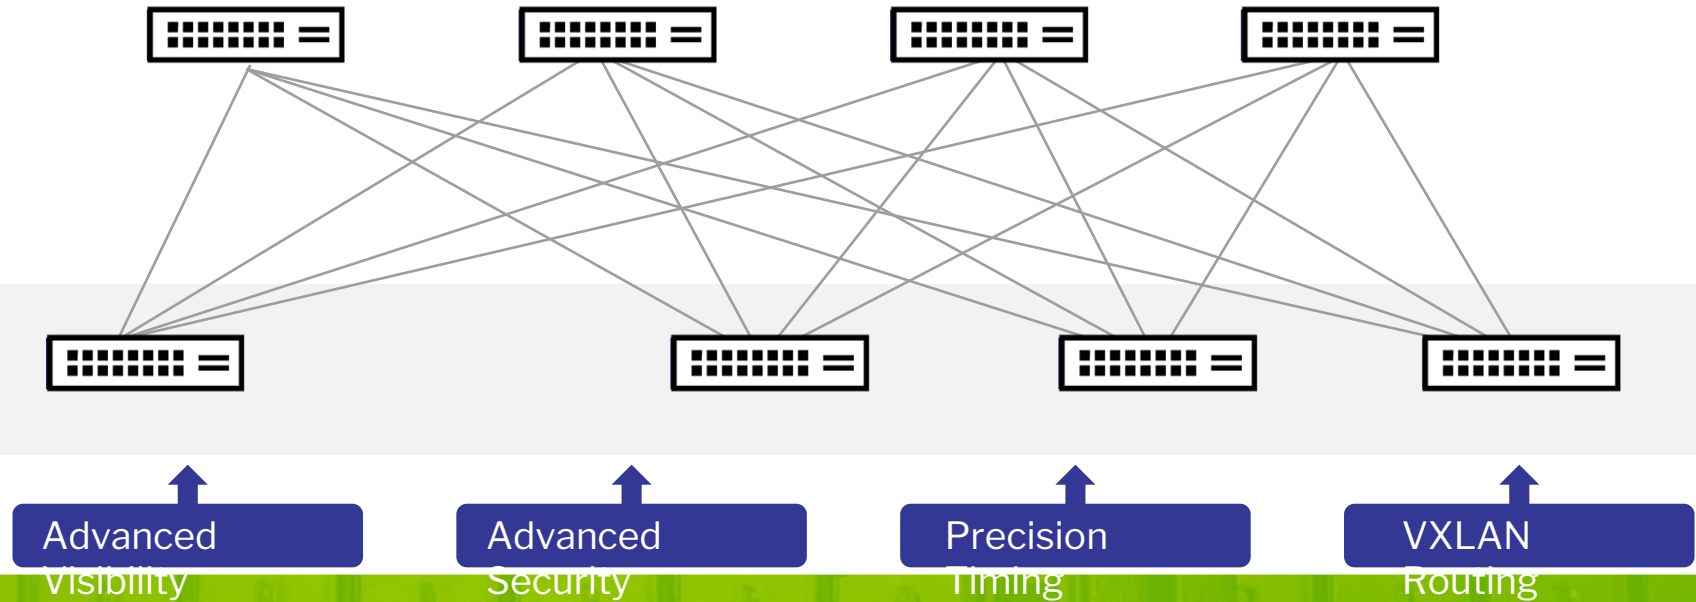
Optimized for Speed

Dumb/Fast Switches

LEAF

Optimized for Features

Smart/Slow Switches




**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.



# Adaptive Routing from NVIDIA



## Traditional ECMP

- Static hashing
- Independent of traffic conditions
- Bigger flows = higher chance for congestion
- High tail latency

## Adaptive Routing

- Congestion based port selection
- Flowlet-aware: eliminates out-of-order Packets
- Multi-vendor friendly
- RoCE OOO placement for highest efficiency

# The Problem

- Compromises were made:

SPIN

Optimized for Speed


Dumb/Fast Switches

LEAF

Optimized for  
Features

Smart/Slow Switches

Adaptive



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# The Problem

- Compromises were made:

SPIN

Optimized for Speed


Dumb/Fast Switches

LEAF

Optimized for Features

Smart/Slow Switches

Adaptive



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.



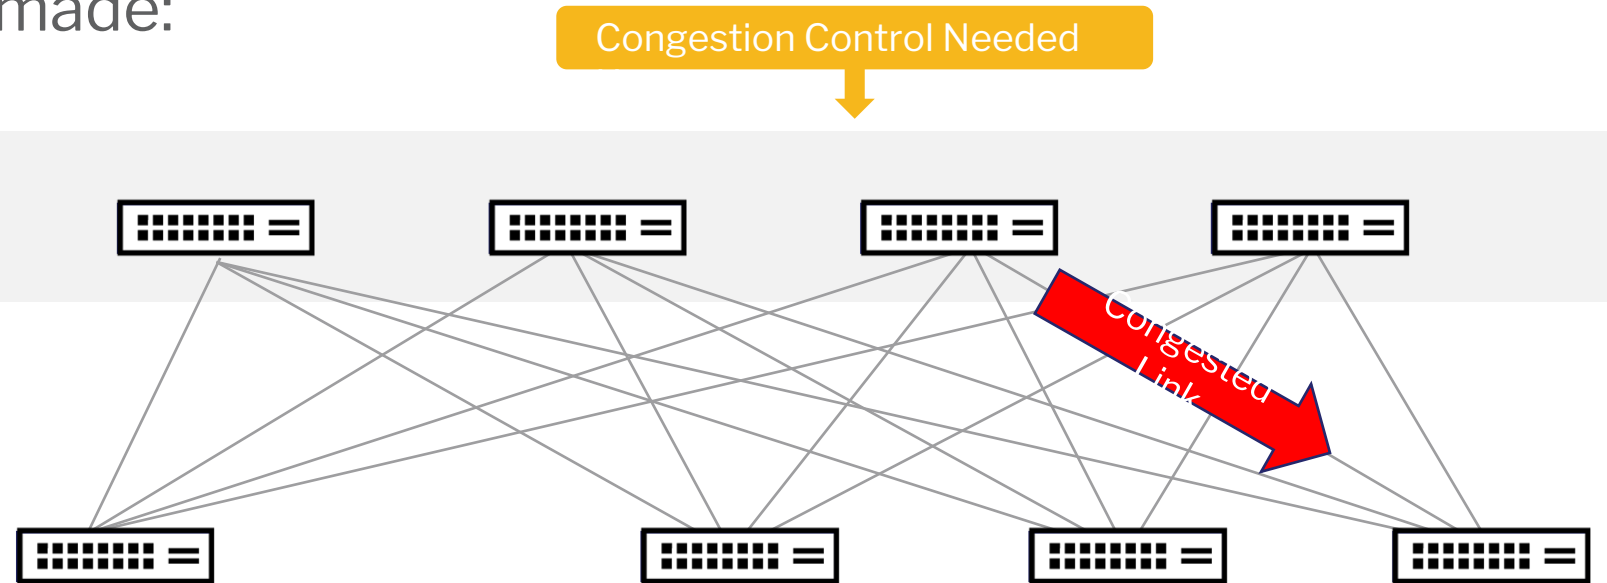
# The Problem

- Compromises were made:

SPIN

Optimized for Speed

Dumb/Fast Switches



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# The Problem


- Compromises were made:

SPIN

Optimized for Speed

Dumb/Fast Switches

Must Never  
Connect to  
Servers



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# Return of End of Row (EoR) Switches

- The world is changing


SPIN

E

LEAF

F

ToR with 4-8 servers per rack?



**OPEN**  
Compute  
Project®


Connect. Collaborate.  
Accelerate.

# Return of End of Row (EoR) Switches

- The world is changing

51Tbps EoR Switch  
100+ Servers Per  
Row

No ToR's Needed




**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# Spectrum-4 SN5600

64x800GbE · 128x400GbE · 256x200GbE

Switch ASIC	<ul style="list-style-type: none"><li>NVIDIA Spectrum-4</li></ul>
Switching Capacity	<ul style="list-style-type: none"><li>51.2Tbps</li></ul>
Ports	<ul style="list-style-type: none"><li>64 Cages: 800G or 2x 400G</li></ul>
System CPU	<ul style="list-style-type: none"><li>x86, Six-Core Xeon</li><li>RAM: DDR4 SDRAM 32GB</li><li>Image storage: SATA SSD 256GB</li></ul>
System Power	<ul style="list-style-type: none"><li>PS: AC, 1+1 redundancy, hot swap</li></ul>
Mounting Options	<ul style="list-style-type: none"><li>Fixed</li><li>Tool-less, Rack mobility</li></ul>
Dimensions	<ul style="list-style-type: none"><li>H: 2U, 3.43'' (87mm)</li><li>W: 16.8'' (428mm)</li><li>D: 26'' (660mm)</li></ul>
Airflow	<ul style="list-style-type: none"><li>N+1 fans, hot swap, forward and reverse</li></ul>



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.



# The Solution: Smarter Networking at Every Tier

- Advanced functionality anywhere you want
  - VXLAN routing
  - Security
  - Visibility
  - Adaptive Routing
  - Congestion Control



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# Thank You!

Connect. Collaborate. Accelerate.




**OPEN**  
Compute  
Project®



# The Problem

- Compromises were made



**OPEN**  
Compute  
Project®


Connect. Collaborate.  
Accelerate.

# The Problem

- Compromises were made

Dumb/fast Switches

Adaptive Routing




**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.

# The Problem

- Compromises were made

Congestion Control needed here



**OPEN**  
Compute  
Project®


Connect. Collaborate.  
Accelerate.



# The Problem


- Compromises were made

Dumb/fast Switches  
\*Must Never Connect to  
Servers



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.



**OPEN**  
Compute  
Project®


Connect. Collaborate.  
Accelerate.

# Return of End of Row (EoR)

- The world is changing

Spine

ToR with 4-8 servers per rack?



**OPEN**  
Compute  
Project®

Connect. Collaborate.  
Accelerate.