Priority Groups
More discussion

Manoj Wadekar

January 2008

Contributors

- Parag Bhide, Emulex
- Craig Carlson, Qlogic
- Claudio Desanti, Cisco
- Dinesh Dutt, Cisco
- Uri Elzur, Broadcom
- Anoop Ghanwani, Brocade
- Bruce Klemin, Qlogic
- Mike Ko, IBM
- Joe Pelissier, Cisco
- Renato Recio, IBM
- J. R. Rivers, Nuova
- Ravi Shenoy, Emulex
- Pat Thaler, Broadcom

Agenda

- Requirements – re-emphasized
- Configuration Tables
- Template config tables
- Summary

DCB: Converged Link

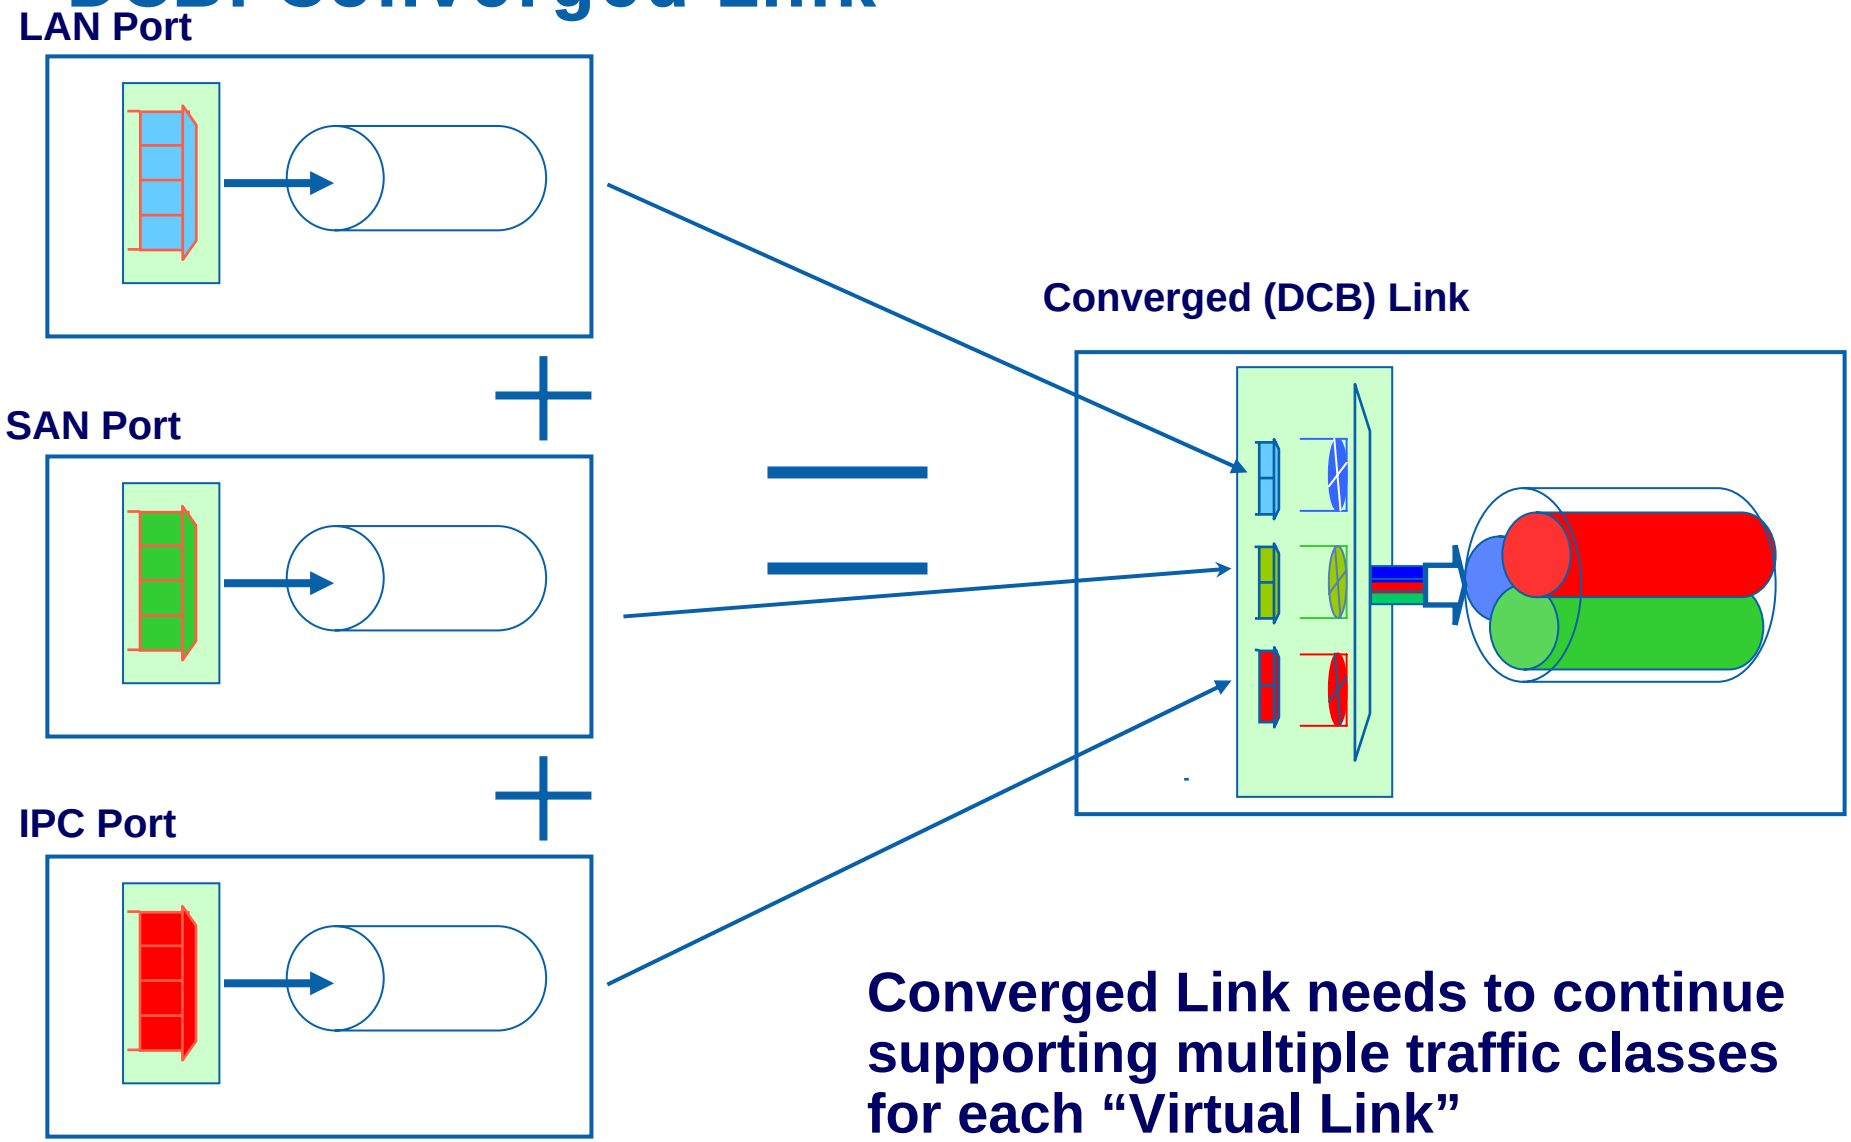

Converged Link needs to continue supporting multiple traffic classes for each "Virtual Link"

Management view of converged link

DCB Cloud has multiple devices that support converged links
- Provide consistent management hooks
- Need to support configuration of BW on DCB link

Configuration for BW assignment for each “Priority Group”
- Example: 40% LAN, 40% SAN, 20% IPC

Should allow multiple traffic classes within “Priority Group”
- Allows these traffic classes to share BW without hard configuration
- Example: VoIP and Bulk traffic to share 40% LAN BW

Can not compromise low latency application due to convergence
- MUST allow strict, high priority scheduling of IPC (and equivalent) traffic

Should provide management infrastructure (MIBs)
- Defining scheduling algorithms is too restrictive and not necessary
- Interoperability for management is important

Term Definition

Pri: Priority

- This is actual marking of traffic on the wire (802.1p bits)

Priority Group (PG) - PGID

- E.g. LAN, SAN, IPC, Management etc.

PG%

- % of Link Bandwidth allocated for a particular PGID

SP (Strict Priority)

- No BW check for this priority – follows strict priority scheduling

Traffic type, queues, scheduling, BW

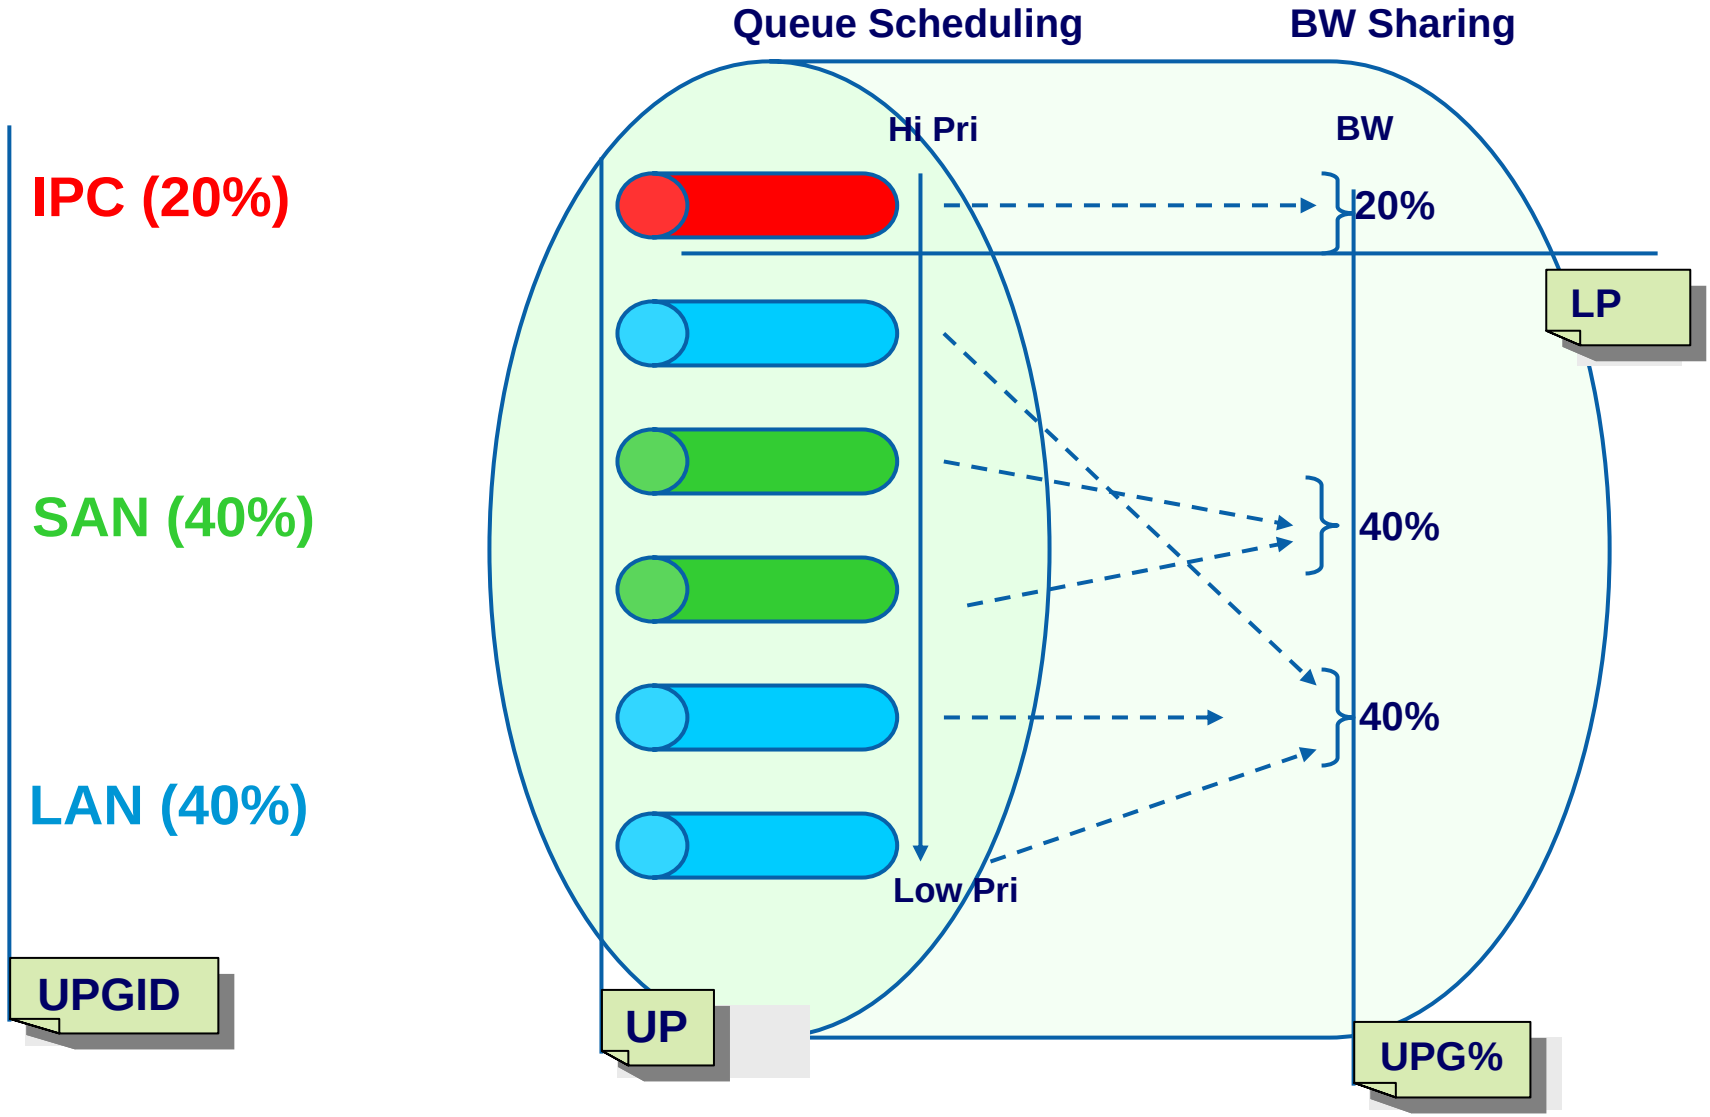

Configuration Tables:

Pri	PGID	Desc
0	2	LAN
1	2	LAN
2	1	SAN
3	1	SAN
4	2	LAN
5	2	LAN
6	NC	NC
7	0	IPC

Table 1: UP-UPGID Table

PGID	PG%	SP	DESCRIPTION
0	-	TRUE	IPC
1	50	FALSE	SAN
2	50	FALSE	LAN
-			
-			
-			
-			

Table 2: UPG-BW Table

NOTE 1: PG% defines MAX BW allowed, if link is fully occupied

NOTE 2: If multiple priorities are mapped to the same group, similar treatment (e.g. no-drop) is expected.

Summary

Allow BW configuration for Traffic Classes

Consistent configuration mechanisms across devices

Maintain low latency treatment of certain traffic classes

Allow configuration of converged link to support BW sharing

Maintain flexibility of implementation algorithms