MultiChannel Discovery and Configuration

v6

IEEE 802.1Qbg, Orlando

March 15, 2010

Contributing Authors

Company	Contacts
BNT	Daya Kamath
BNT	Jay Kidambi
BNT	Vijoy Pandey
Broadcom	Uri Elzur
Brocade	Anoop Ghanwani
Emulex	Chait Tumuluri
HP	Paul Bottroff
HP	Paul Congdon
HP	Chuck Hudson
HP	Michael Krause
IBM	Vivek Kashyap
IBM	Renato Recio
IBM	Rakesh Sharma
Juniper	Srikanth Kilaru
QLogic	Manoj Wadekar

Multi-Channel: Multiple isolated channels (S-VLANs) sharing a single physical LAN.

Provides the ability to support VEBs, VEPAs, and individual links on a single physical link.

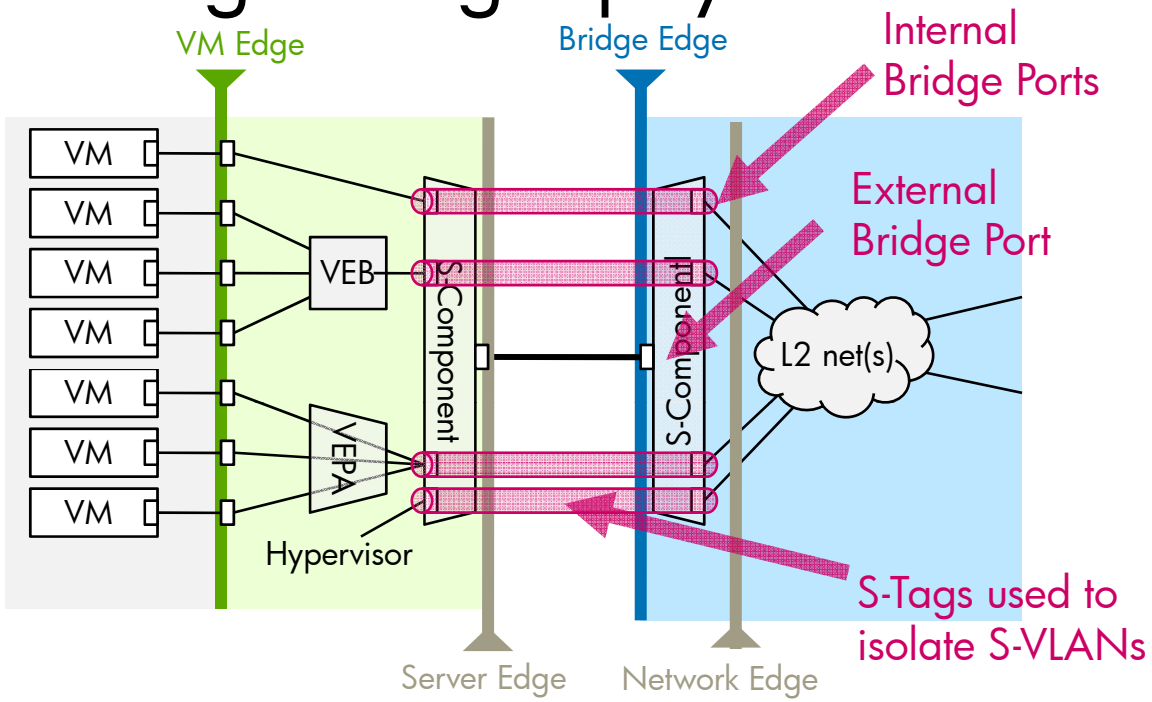

- Channels (or S-VLANs) are implemented using a reduced function S-VLAN aware Bridge relay in both the station and bridge. (proposed re-naming to an E-Comp to indicate the use in a MultiChannel application)
- The S-Comp has two port types. These are:
  - A single externally facing Bridge Port (also can be a 802.1ag aggregation group) which is a specialized Provider Network Port (PNP) conforming to the 802.1Qbc definition for Remote Customer Access Port. (proposed naming a Server to Bridge Access Port (SBAP))
  - Many internal Bridge Ports (previously called vPorts, one for each S-VLAN) which are specialized Customer Network Ports each coupled in a 1-1 relationship by internal LAN to a Bridge Port of a C-VLAN aware Bridge relay or ES (VEPA, VEB or ES in station or C-Comp or VLAN unaware in Bridge). (proposed naming a Network Access Port (NAP))

MultiChannel Discovery and Configuration

- MultiChannel Discovery Protocol (MDP) runs over LLDP at the LAN level to establish S-VLANs
- The MDP protocol allows the station to request S-VLAN configurations and the Bridge to supply (or deny) its best fit configuration to the station.
- The configuration of VEPA, VEB or ES attached to an S-VLAN is handled by other protocols executed after S-VLAN establishment:
  - Edge Virtual Bridge Discovery Protocol (EDP) runs over LLDP at the S-VLAN level and is used to establish the protocols running over an S-VLAN (channel).
  - Virtual Station Interface Discovery Protocol (VDP) runs over Edge TLV Transport Protocol (ETTP) over an S-VLAN and is used to establish VSI for attachment of VMs.

Bridge Architecture for MultiChannel

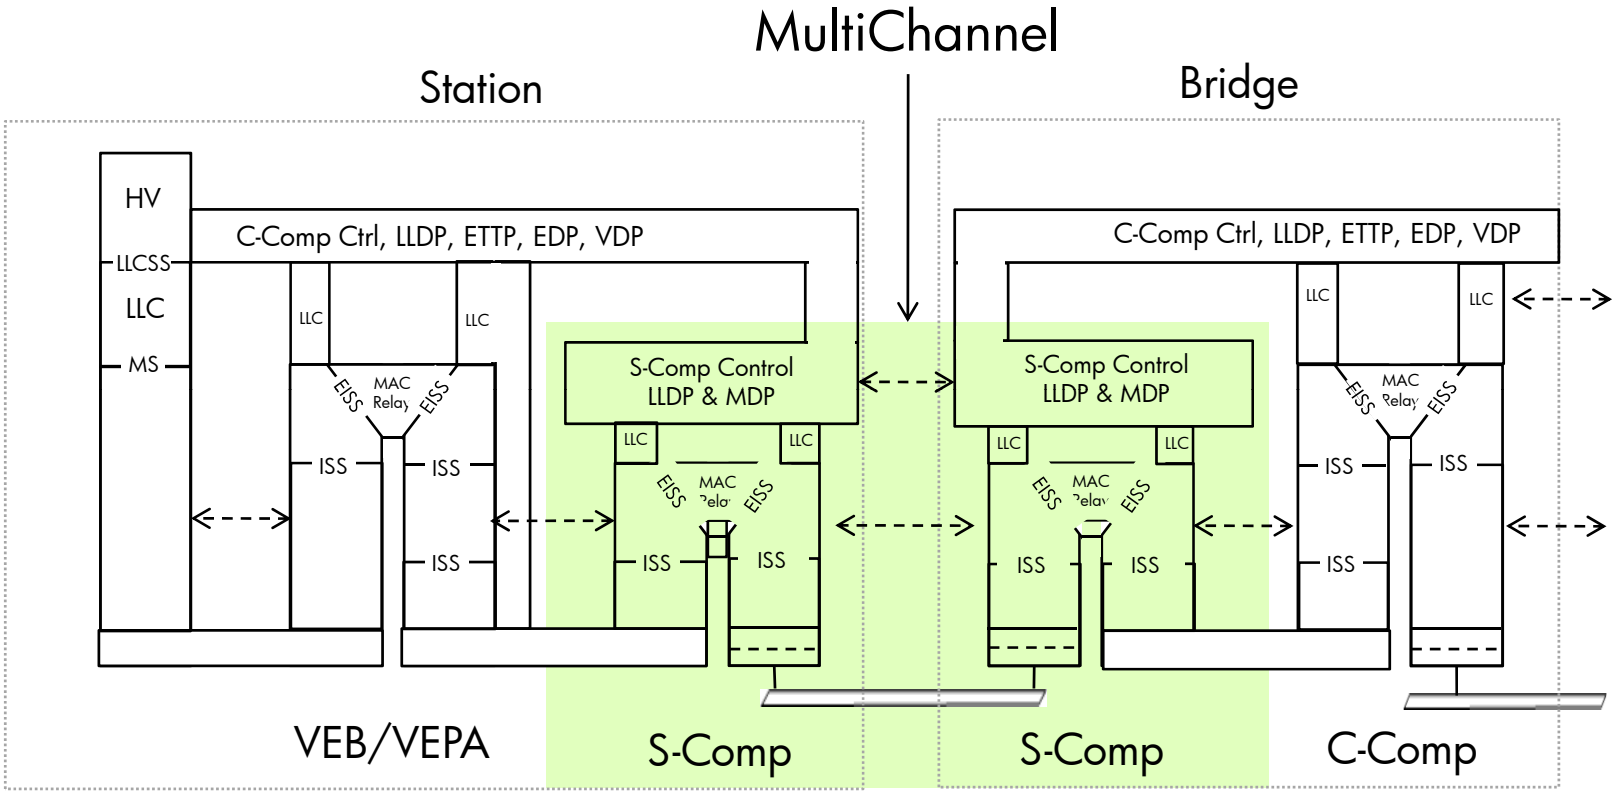

- The S-Comp used to implement MultiChannel has reduced function likely conforming to the 802.1Qbc Port-mapping S-VLAN component

Some Terminology

- **S-VLAN**: a channel
- **S-Component**: the S-VLAN aware Bridge Component (from 802.1Qbc Port-mapping S-VLAN component, proposed new name E-Comp)
- **RCAP**: a Remote Customer Access Port of an S-Comp is an externally facing Bridge Port connecting to the LAN spanning between a station and Bridge (proposed new name SBAP)
- **PAP**: a Provider Access Port of an S-Comp is an internal Bridge Port which connects 1-1 to a VEB/VEPA, C-Comp, VLAN unaware component, or ES (proposed new name NAP)
- **MDP**: an LLDP TLV exchange protocol used to discover and configure E-Components and S-VLANs
- **EDP**: an LLDP TLV exchange protocol used attach a VEPA, VEB or ES to an S-VLAN
- **VDP**: a protocol used over an S-VLAN to establish Virtual Station Interfaces (ES Interfaces through a VEPA, VEB or direct S-VLAN) for VM attachment.

MDP Runs Over LLDP

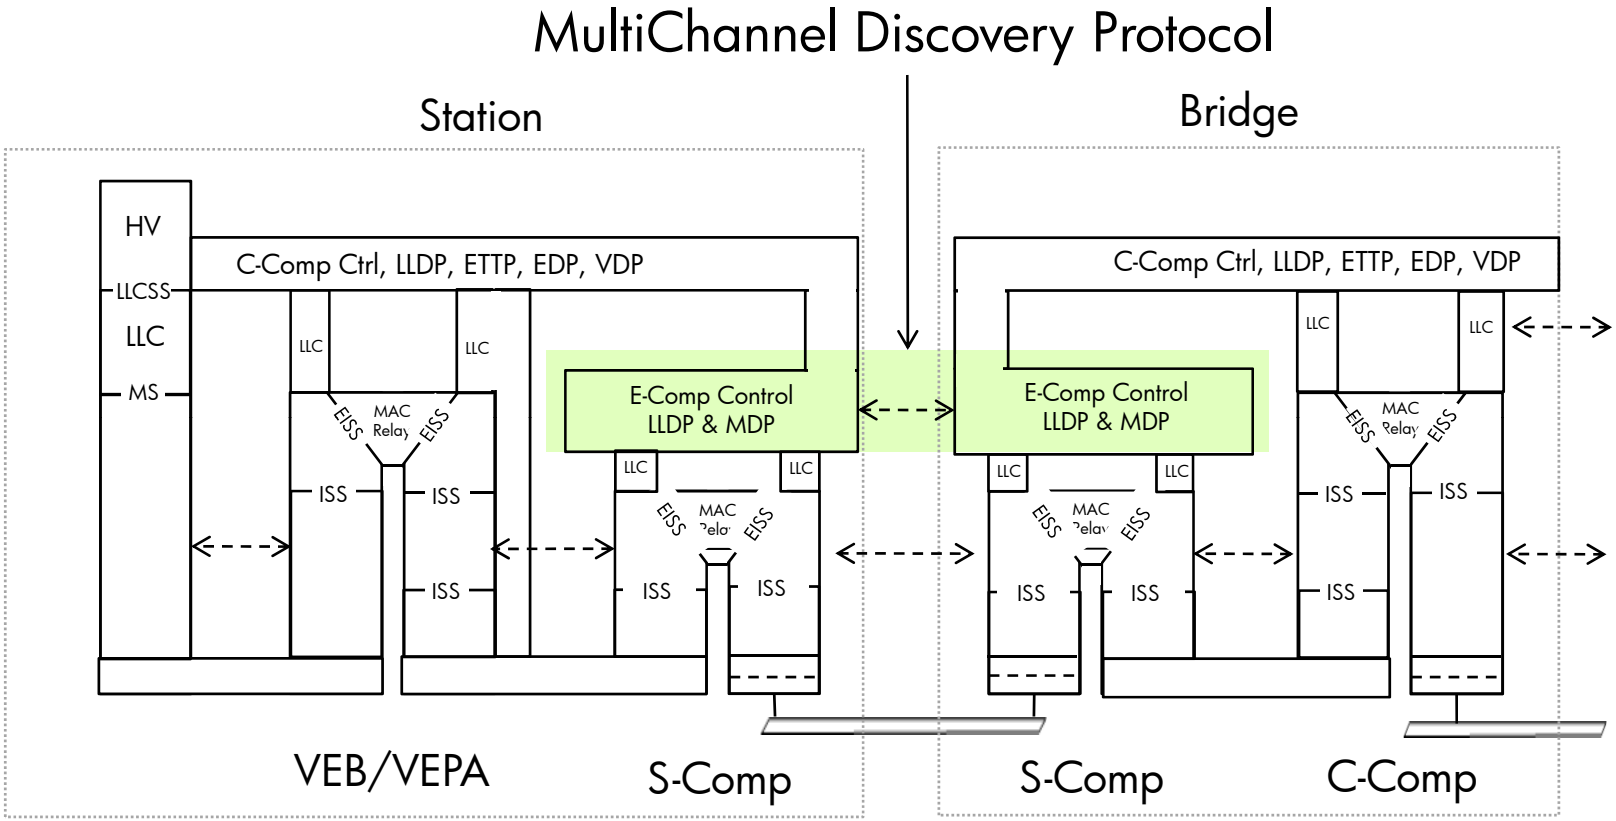

- MDP is a Bridge Control Protocol run in the S-Component control plane.

MDP TLV

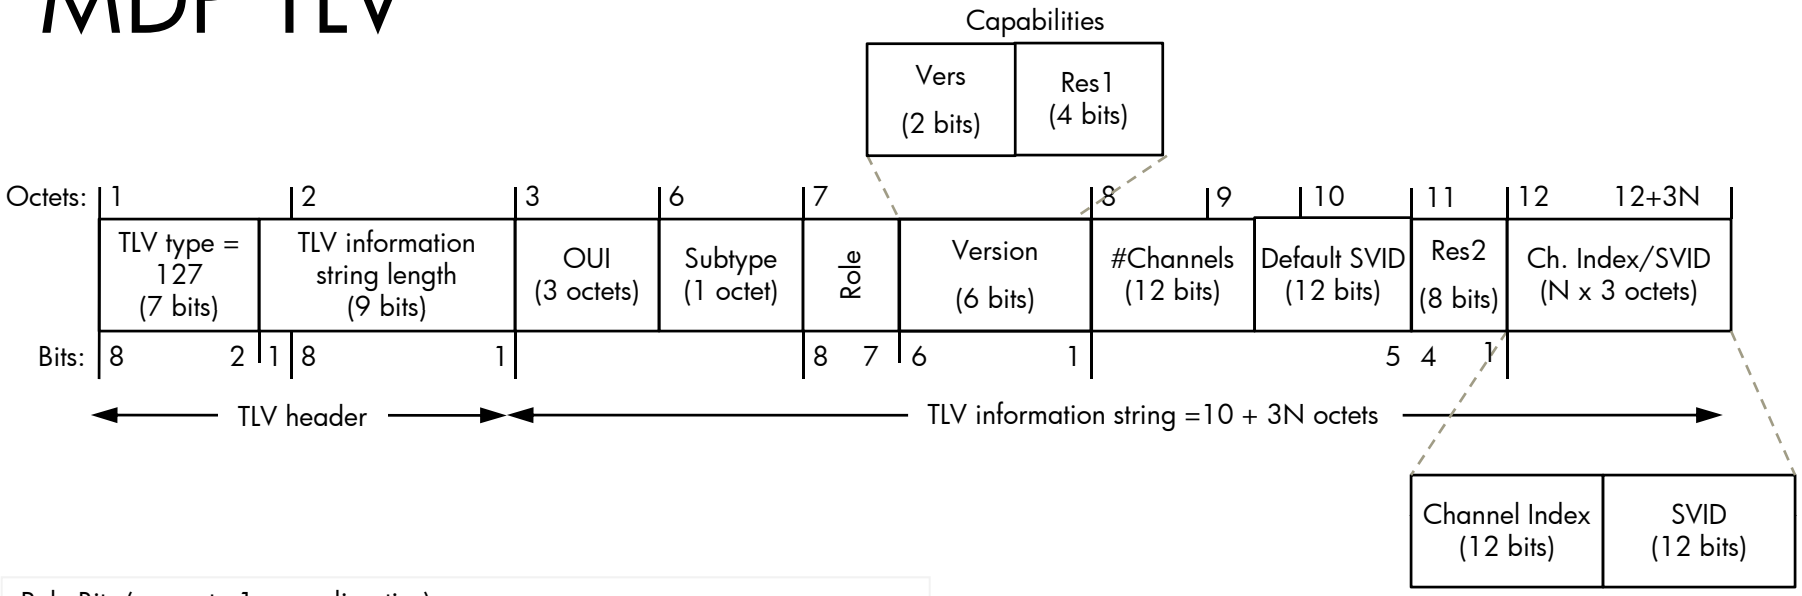

Role Bits (see note 1 regarding ties)

S(01b) – Indicates the sender assigns channels numbers and a default SVID for the default channel 1 and requests SVID assignments from the neighboring 'B'.

B(10b) – Indicates the sender accepts multichannel configuration requests from its neighboring 'S' and that the sender will do the best it can to fill the SVID assignment requests from the neighboring 'S'.

- Version– Describes multichannel capabilities that can be supported by the sender.
  - Vers: 10b identifies this version, 00b disables MCh
  - Res1: must be set to zero, ignored on receipt
- # Channels Supported – Identifies the number of SVID channels that are supportable by the sender.
- Default SVID – Reserved for future use.
- Note1: Channel 1 is the default channel and is always listed first in the list.

- Res2 – must be set to zero, don't run if non-zero
- Ch#/SVID Pairs (for format – 0000b)
  - Channel # -- indicates the index number of the channel. The station assigns channel numbers in the range 0-167. Zero is reserved. Channel number 1 is the default channel and is always the first channel in the list of pairs. The channel index should be between 1 and the maximum number of channels supported by the port.
  - SVID – The S-Tag VLAN ID assigned to the channel. The Bridge assigns SVIDs to channels in the range 1-0xffe. A station uses the 0 SVID to request an SVID assignment from the Bridge.

Note2: A maximum of 167 channels can be supported using format 0000b. Other formats (assuming sequential SVIDs) could be defined to allow support for 2K+ channels.

Note3: This listing could be sparse (in order to indicate arrival and removal of channels). The channel going away is recognized by that channel index/SVID pair is removed.

Note4: Channel 1 is the default channel and is always listed first in the list.

MDP TLV Initial Exchange

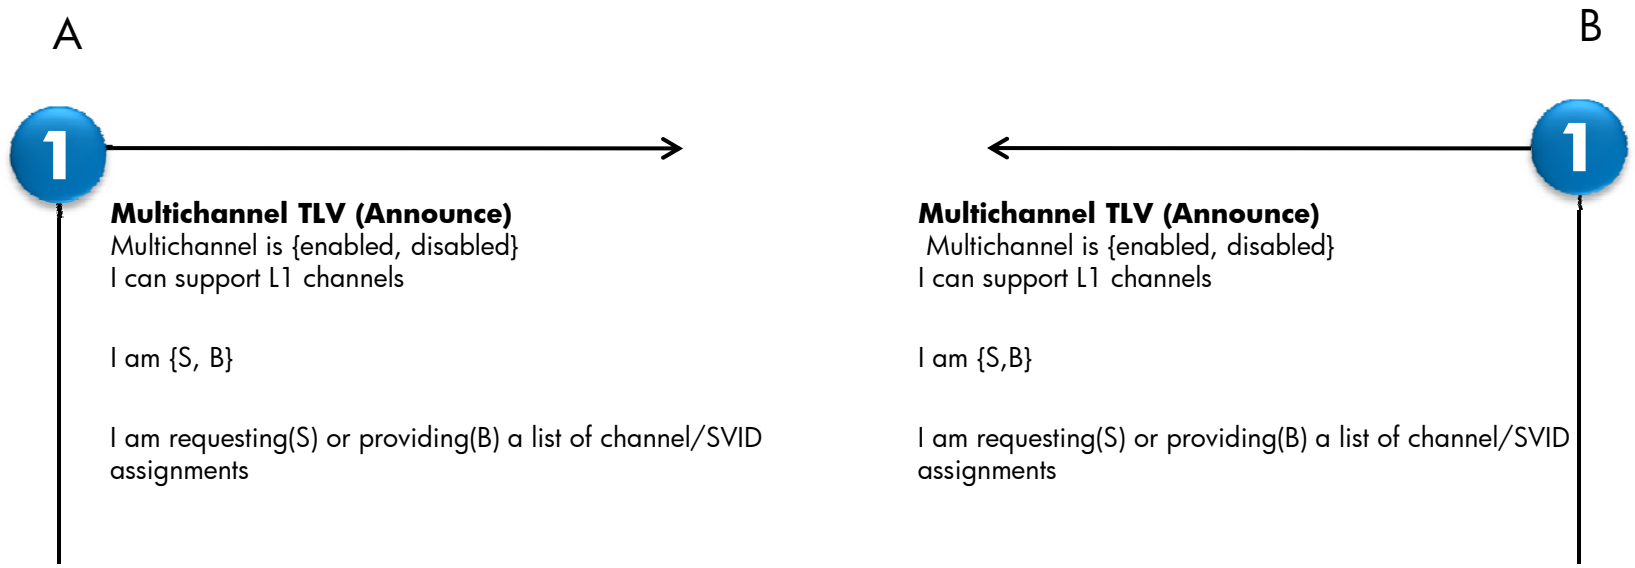

- If both sides are B or both sides are S then MDP will not configure MultiChannel. In the case of two Bridges, then one must take the S role while the other takes the B role.

Some Rules for Channel Numbers

- Capacity is the max number channels supported.
  - Channel numbers start at 1
  - The range for valid channels is from 1-0xffe
  - Channel number 0 may be used to reserve TLV space
  - The S role decides what channel numbers will be active
  - The Bridge always assigns the SVIDs
- Channel requests are made using a Channel/SVID pair where the SVID is 0
- Channel configurations are indicated using a Channel/SVID pair where the SVID is a valid SVID in the range 1-0xFFE.

Basic Success Scenario

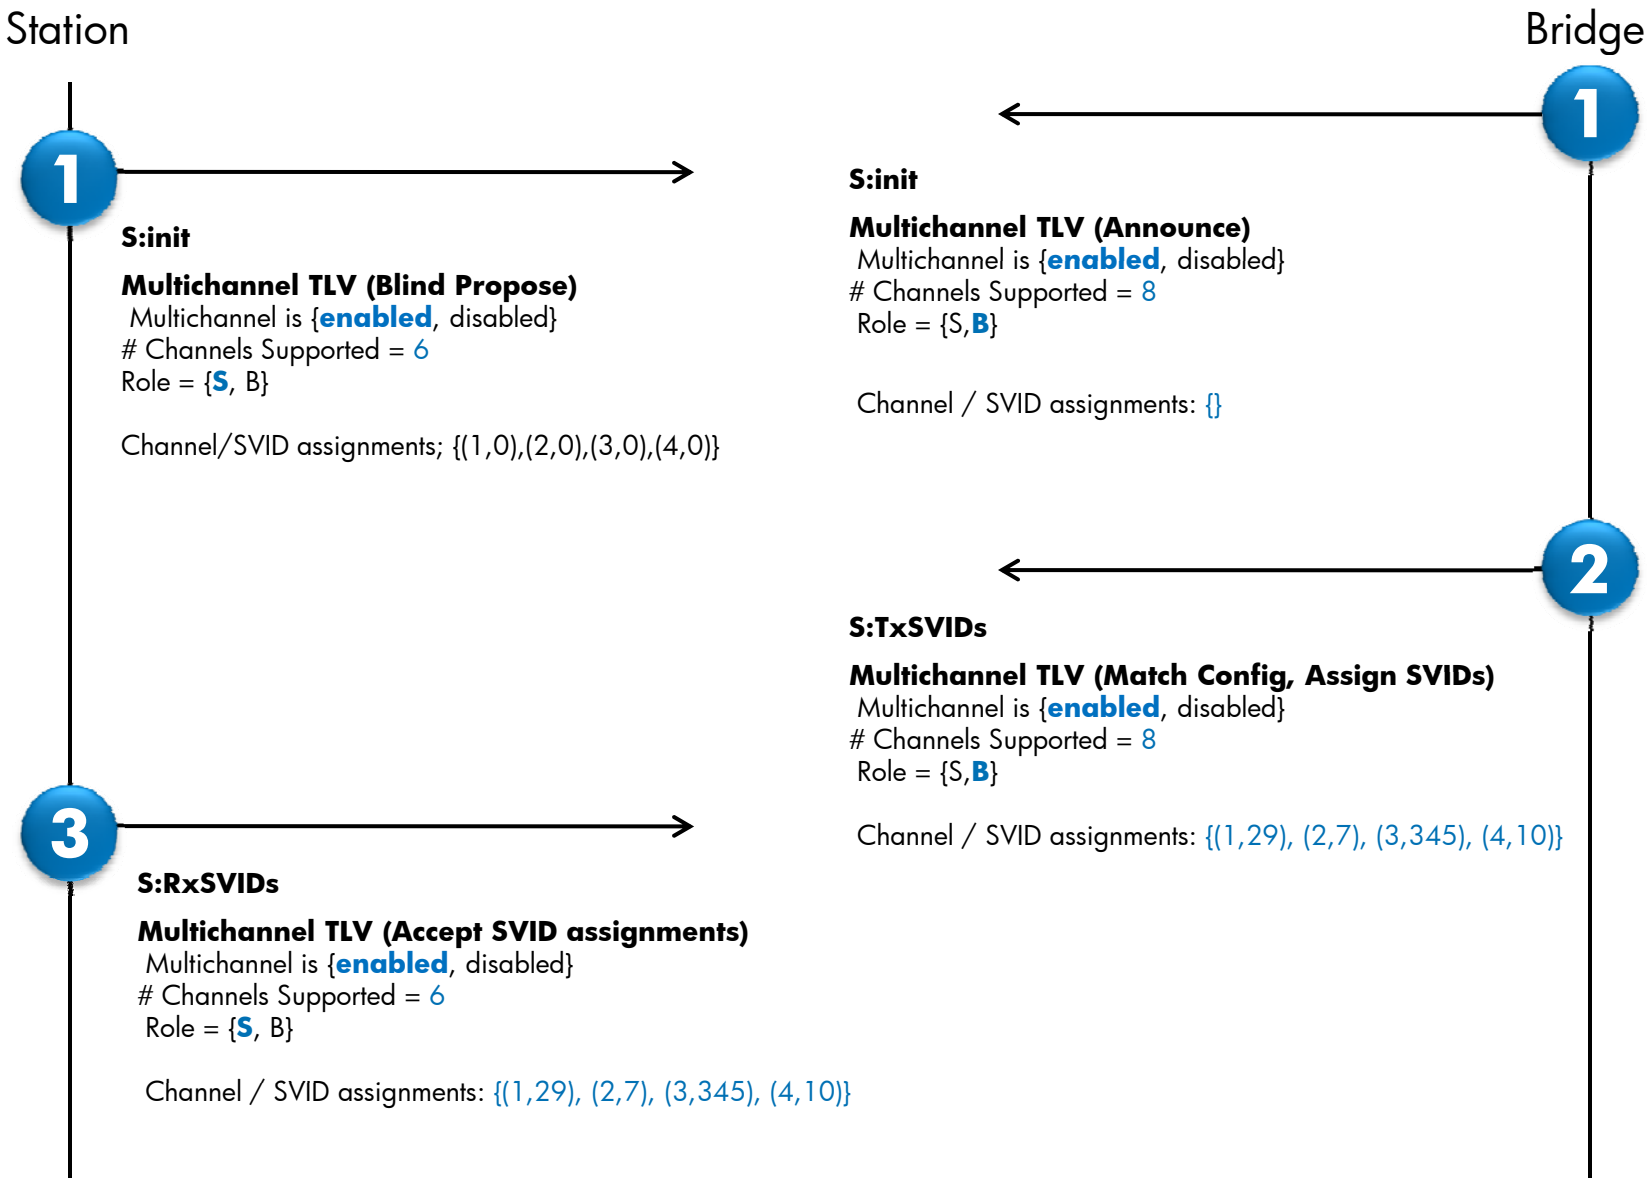

Basic Success Scenario (Sparse)

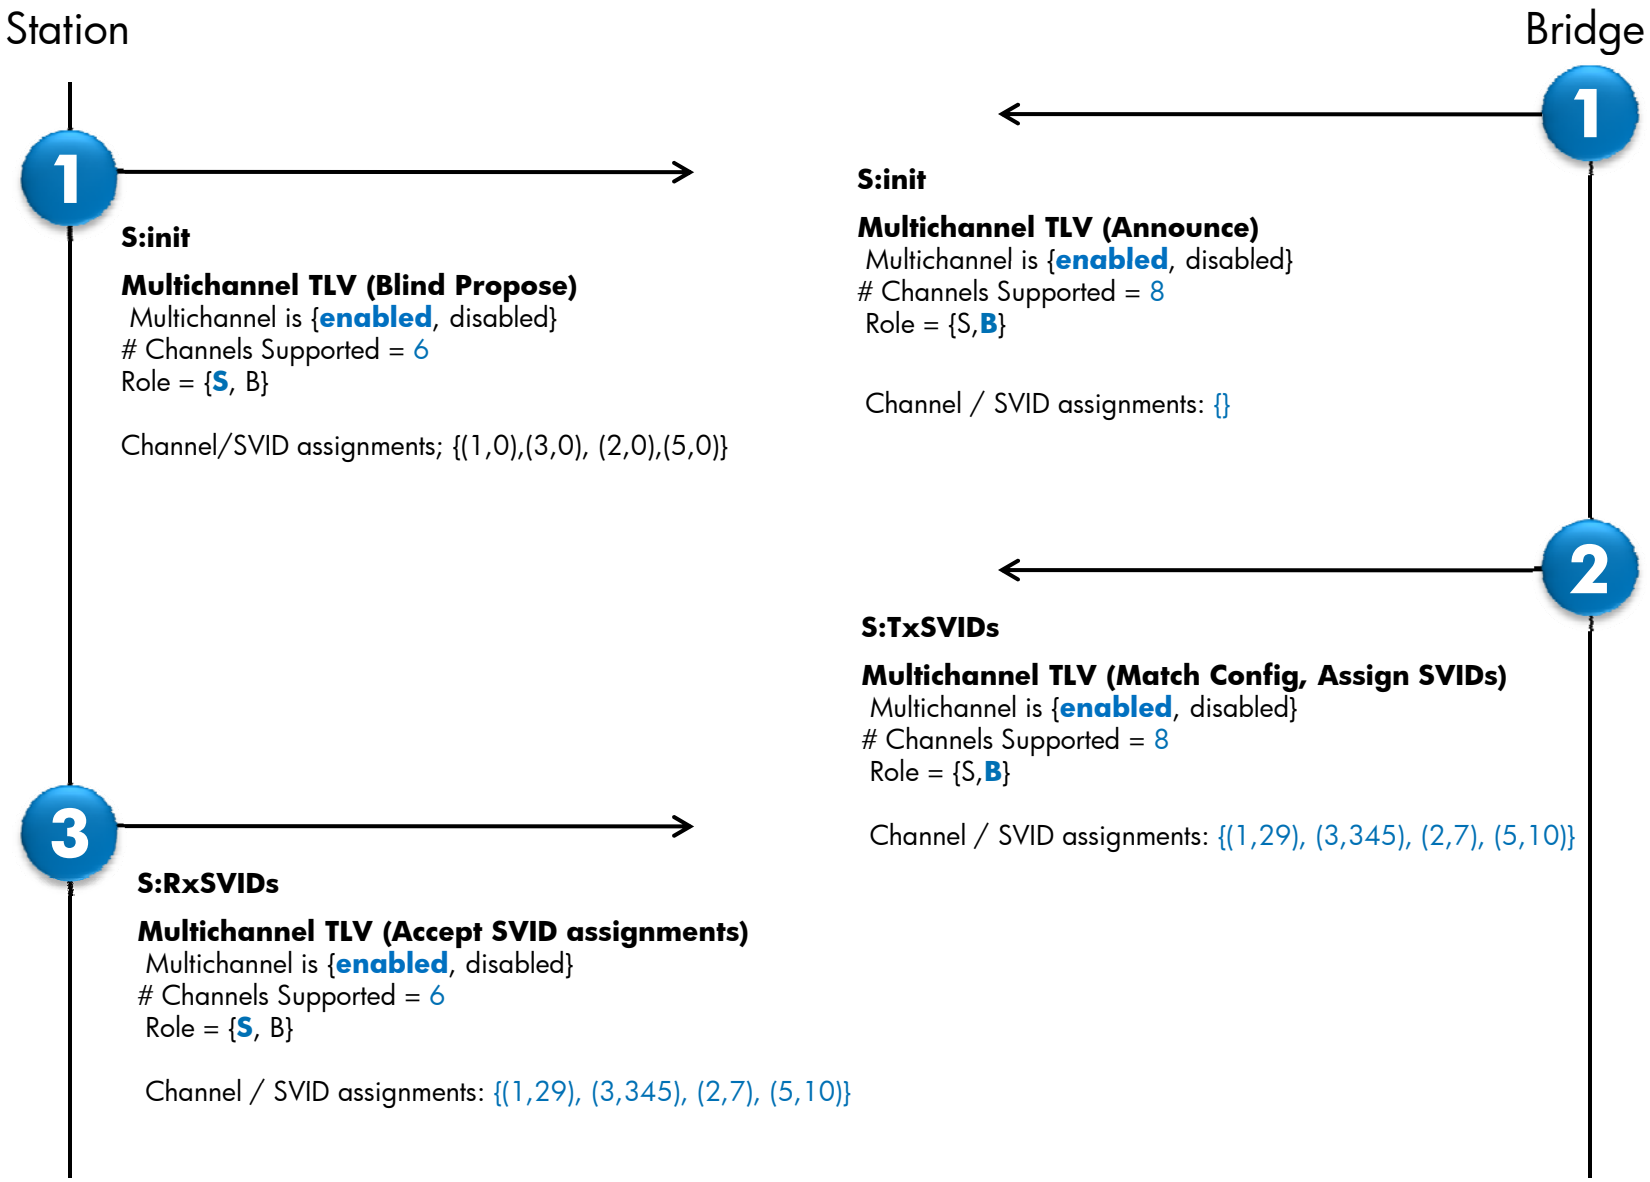

Station adds a channel

continues from basic success scenario

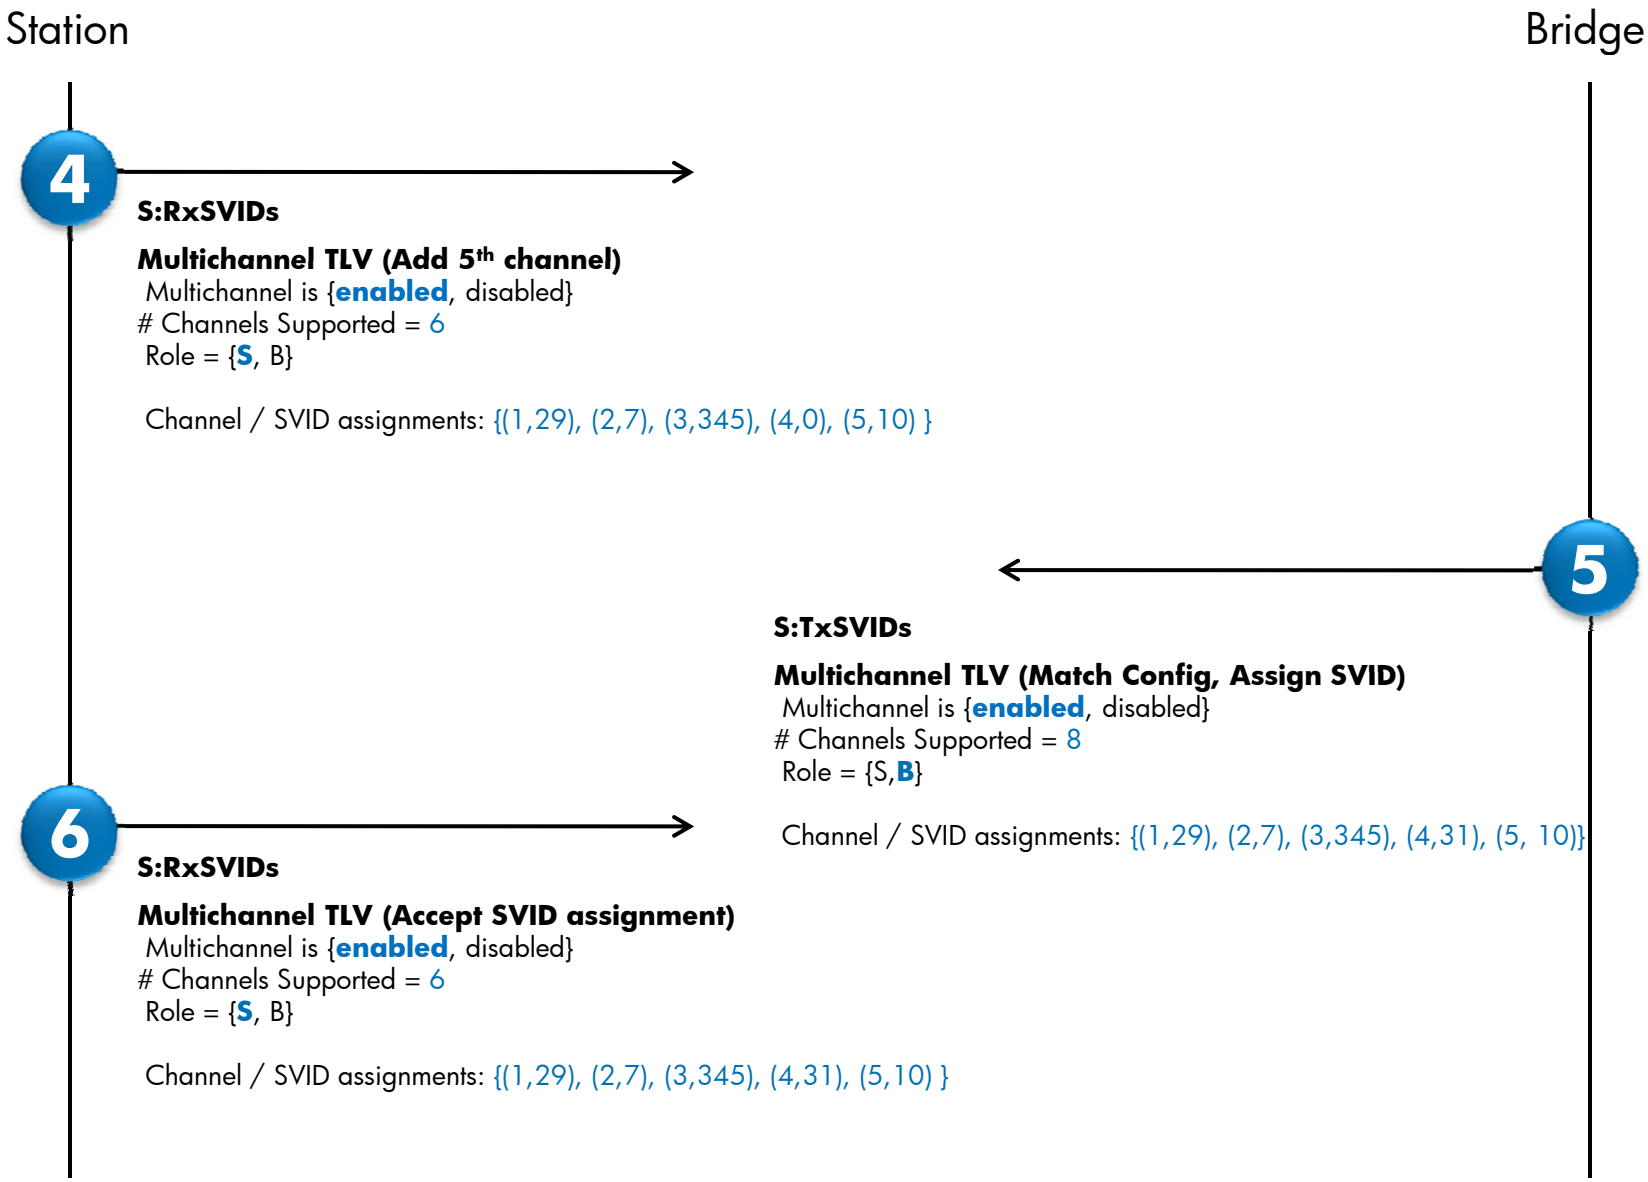

Station Removes a channel

continues from station adds a channel scenario

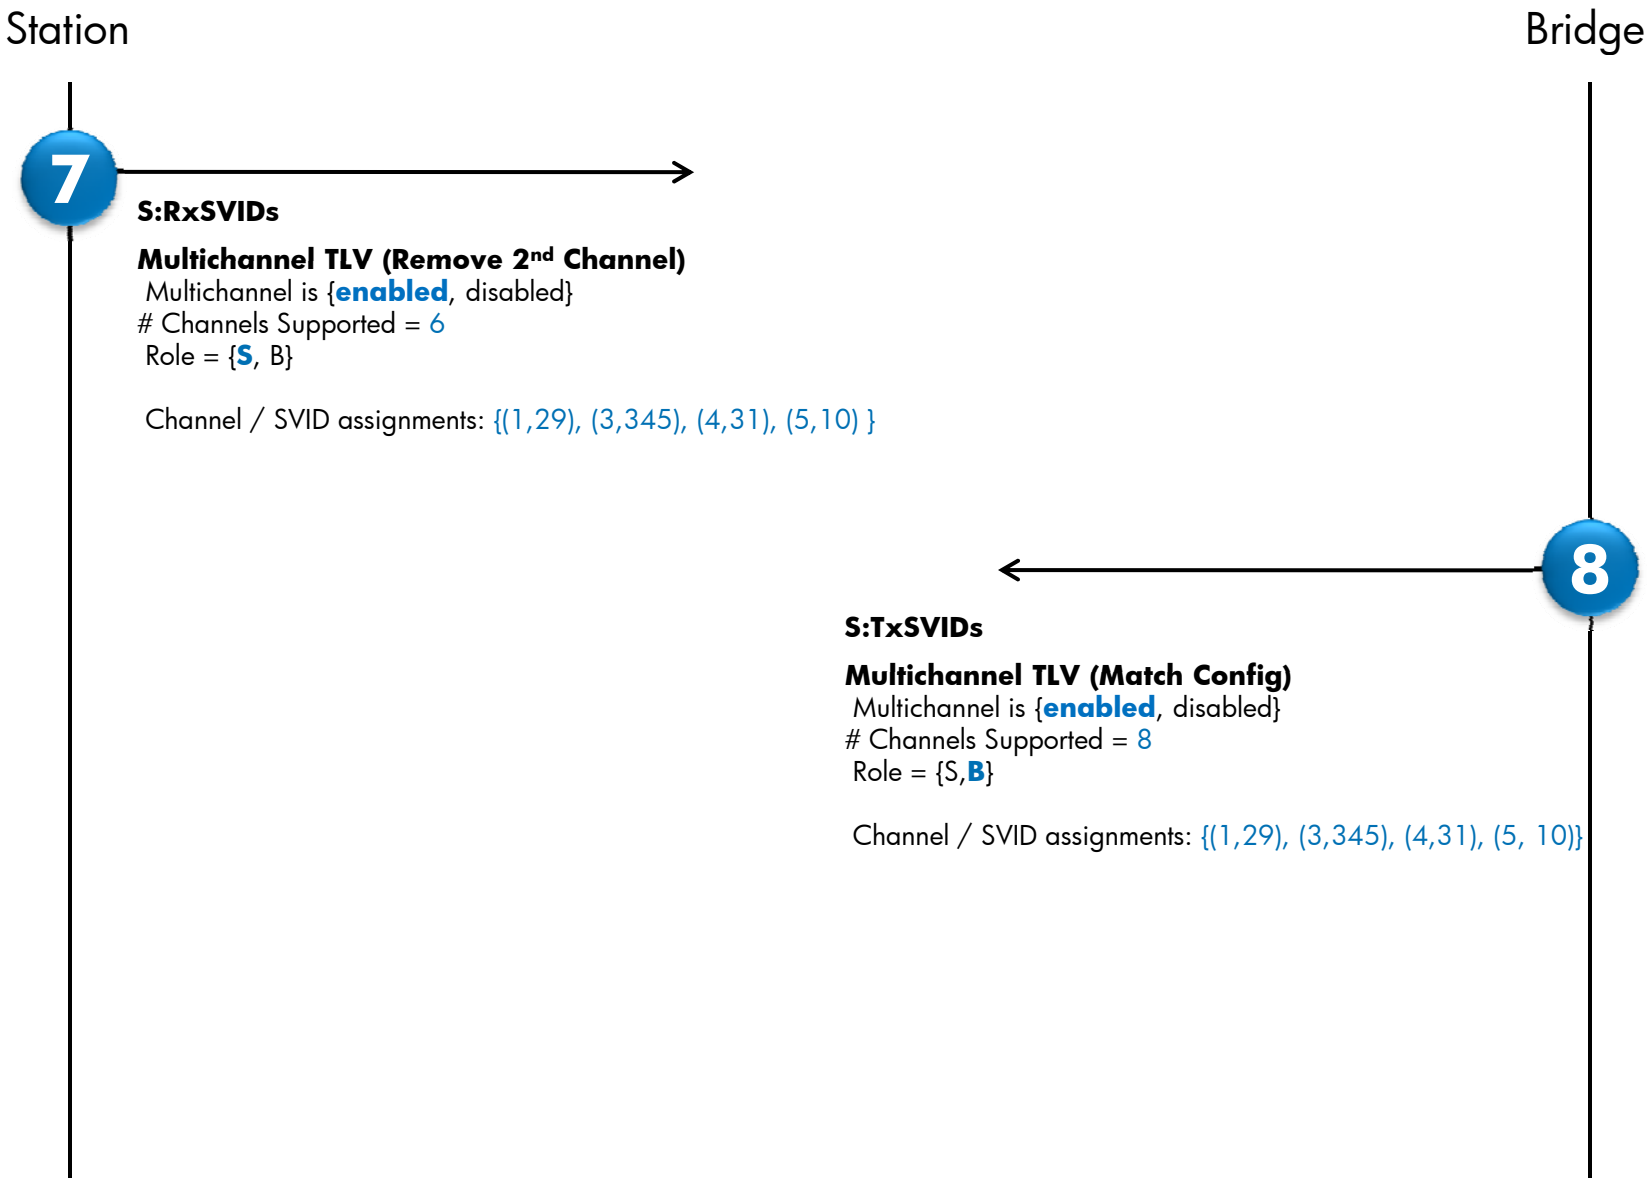

Drop #channels supported

continues from basic success scenario

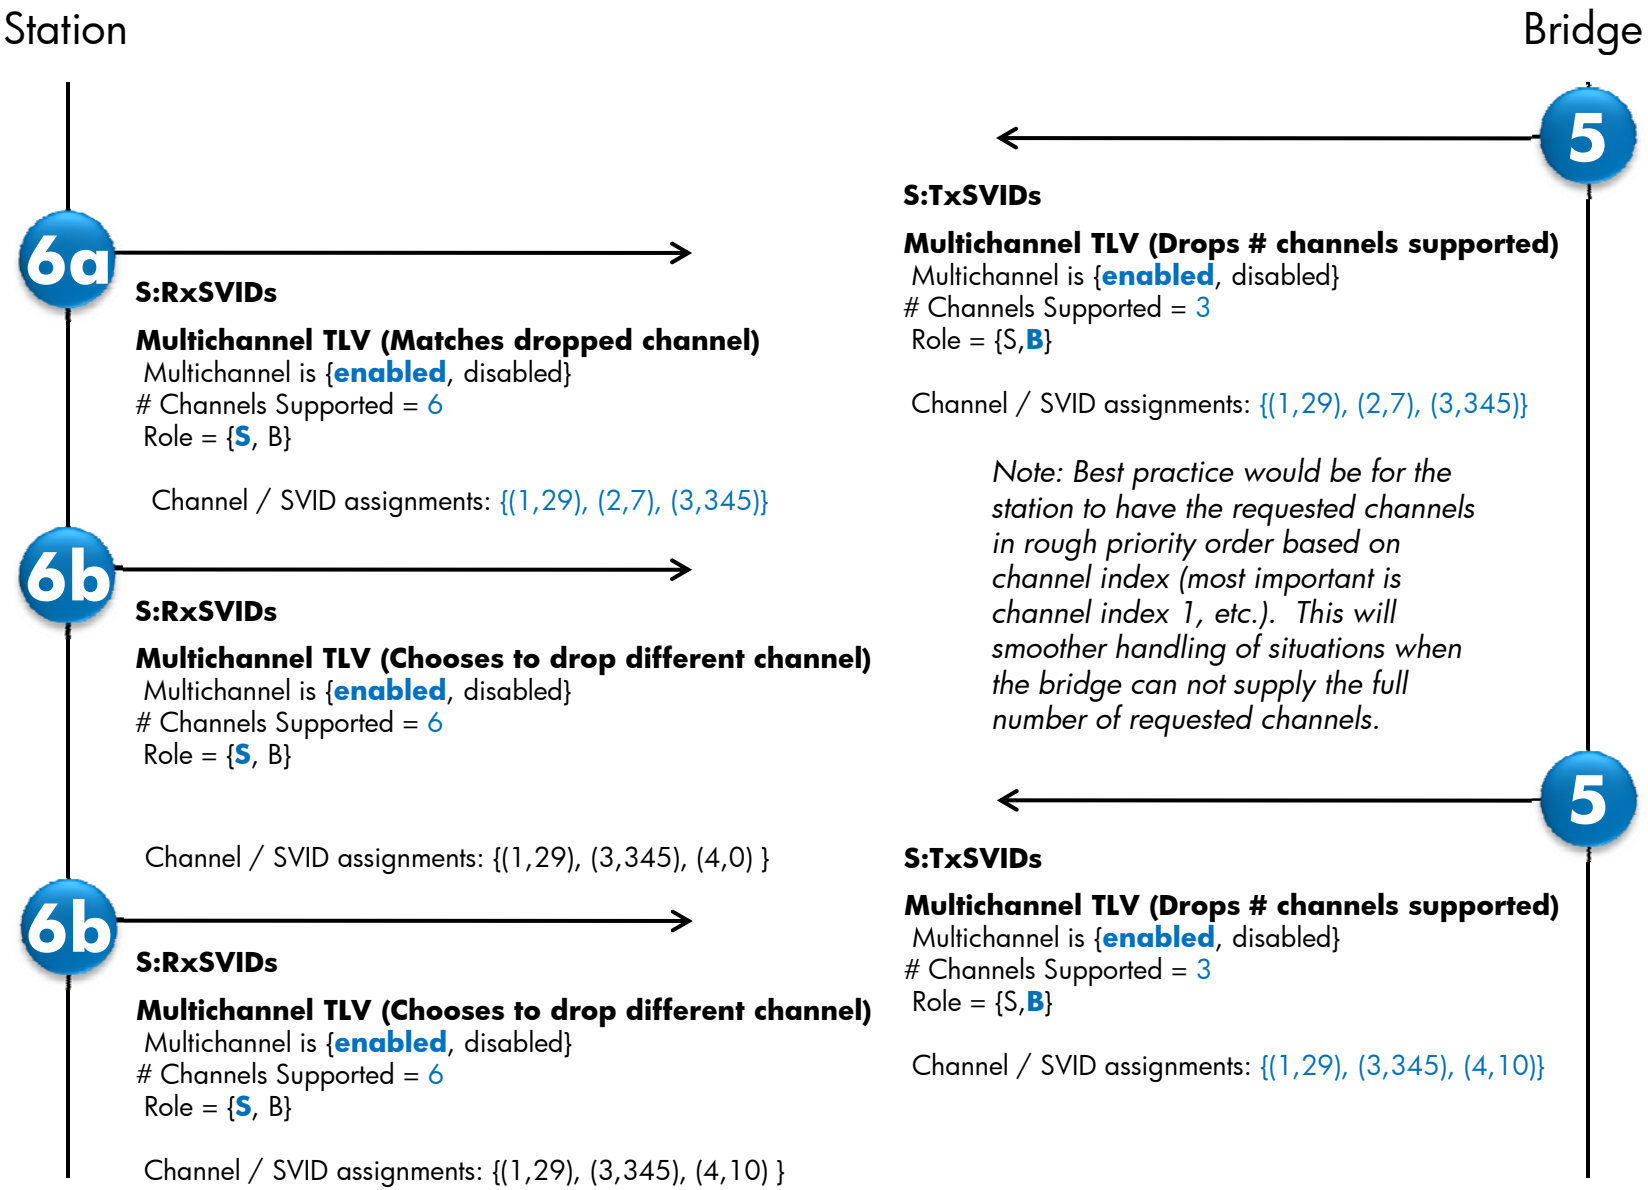

Insufficient channels on bridge

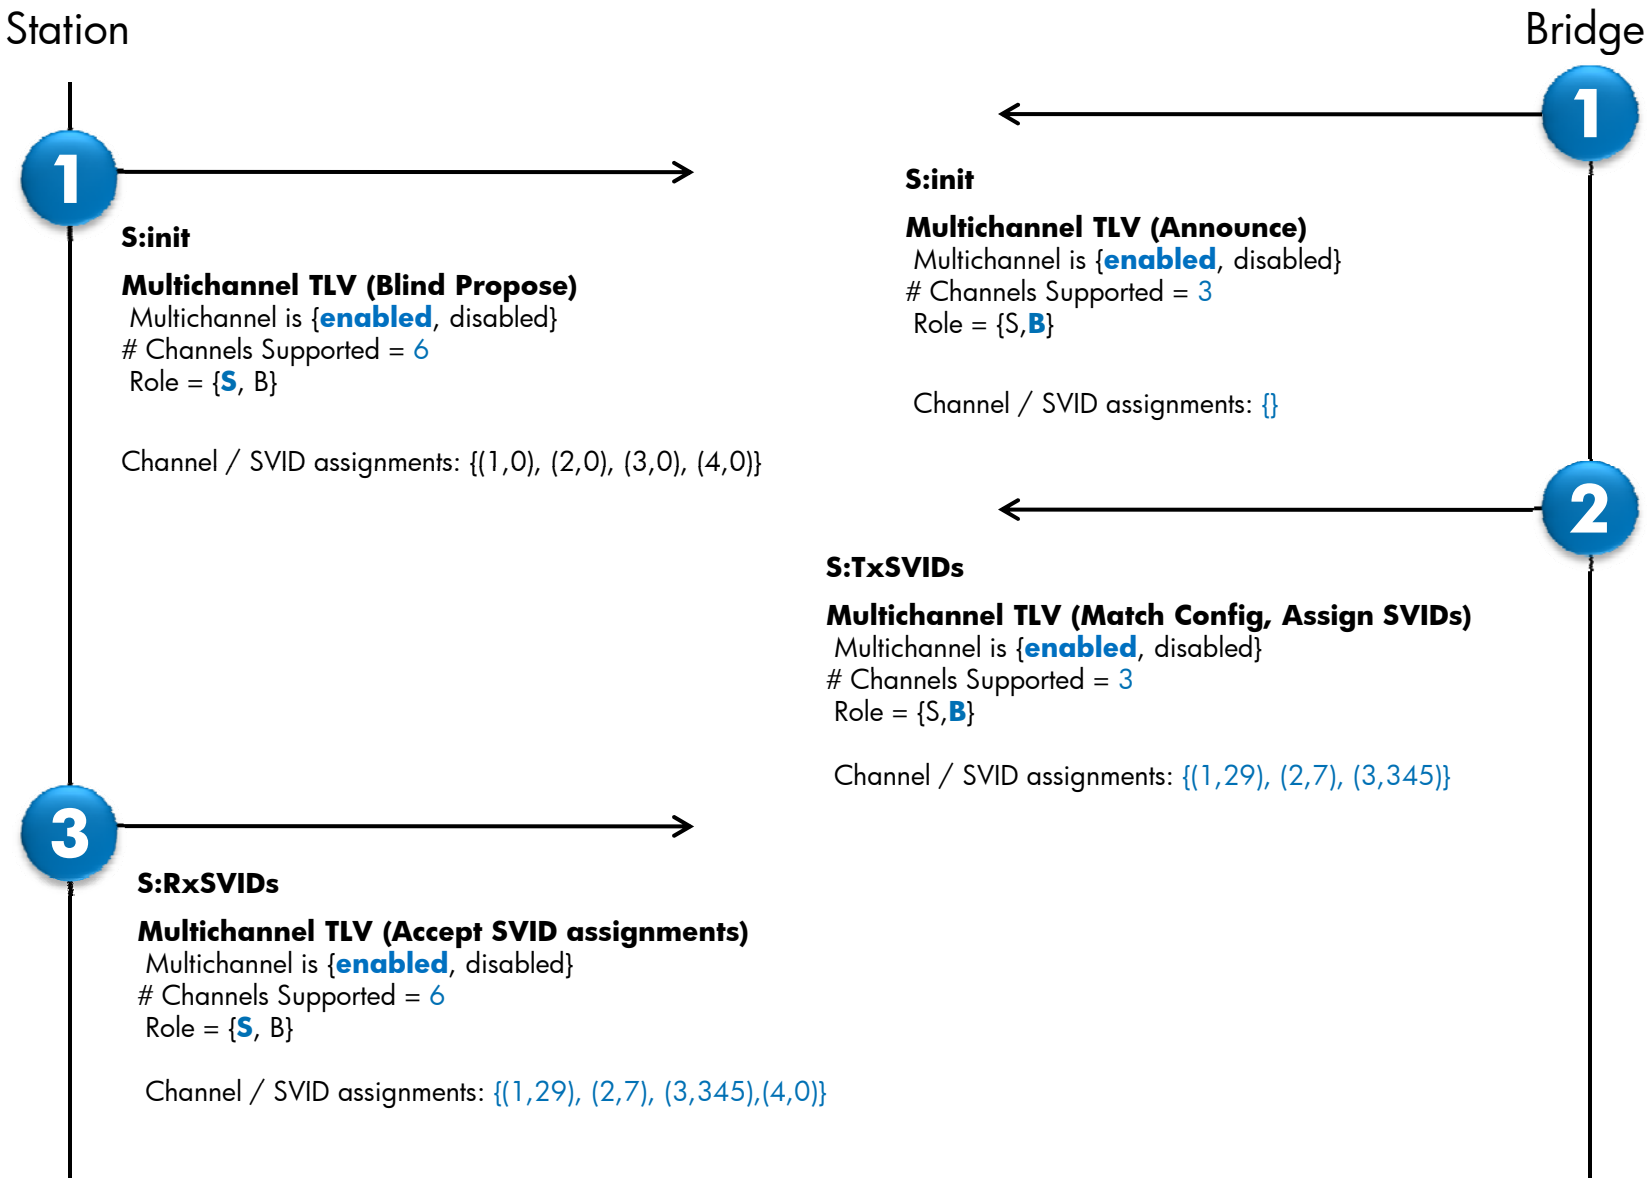

MDP State Machine Variables

- Each equipment (Bridge or Station) participating in the MDP protocol is configured through administrative parameters to enable MultiChannel, to take the S or B role in the dialog, with the channels needed, and with the SVIDs desired or available for configuration.
- The inability to do MultiChannel is reflected by the lack of a MultiChannel TLV in the LLDP database. In this event MultiChannel configuration will not proceed.
- One side must take the S role and the other side must take the B role to allow configuration of MultiChannel. If two Bs or two Ss are connected MultiChannel will not be configured.
- Once we have valid roles the S side may make resource requests and the B side will attempt to fill these requests.
- During operation of MultiChannel the S side may alter its request and the B may change its configuration dynamically.

Variables: Role Bits

- AdminRole: Is the administratively configured value for the local port's role parameter. The value of AdminRole is not reflected in the MCh TLV. The AdminRole may take the value S or B. S indicates the sender is unwilling to accept multichannel configuration (mode, # channels supported, channel index) from its neighbor and that the sender is willing to accept SVID assignments from the neighbor. Stations usually take the S role. B indicates the sender is willing to accept multichannel configuration (mode, # channels supported, channel index) from its neighbor and that the sender is willing do the best it can to fill the SVID assignments from the neighbor. Bridges usually take the B role.
- OperRole: The current operational value of the Role parameter in the local port. This value is included as the Role parameter in the MCh TLV and may take values S or B as described for AdminRole.
- RemoteRole: Indicates the value in the remote MCh TLV role field. rwNull indicates either the TLV was not present in the last LLDP PDU or that no LLDP PDUs have been received. rwS and rwB indicate that the Role field was set in the MCh TLV received and that it had a value of S or B respectively as described for the AdminRole variable.
- mchState: The current running state of MultiChannel. The values this variable can take are NOTRUNNING or RUNNING.

Adminstrative Variables:

- AdminVersion: The administratively configured value for the MCh version parameter. This values is included as the MCh Version parameter in the MCh TLV. If the value is DISABLE = 000b it means MCh is disbled. If the value is VER0 = 100b it means this version.
- AdminChnCap: The administratively configured value for the Number of Channels supported parameter. This values is included as the # Channels supported parameter in the MCh TLV.
- AdminSVIDWants: The administratively configured value for (channel,SVID) pairs wanted by a S. Not used by a B. The value NONE means no channels are wanted. The channel numbers may be any valid number from 0-0xffe. SVID numbers from 1-0xffe mean the S wants this SVID, however will take any SVID assignment the S gets. If the SVID value is 0 it means the S is requesting any available SVID, but has no preferences. This value is used to form the (channel,SVID) pairs in the MCh TLV.
- LastSVIDWants: A local temporary copy of the AdminSVIDWants.
- LocalSVIDPool: The set of SVIDs and bridge ports available for MCh assignment. These are determined by both administrative resource assignments and by resource availability. The OperSVIDList for a B role must be drawn from the LocalSVIDPool.
- LastLocalSVIDPool: A temporary copy of the LocalSVIDPool.

Local Operational Variable

- OperVersion: The current value for the MCh version parameter. This value is included as the MCh version parameter in the local MCh TLV. The value VER0 = 100b means this version. The value DISABLE = 000b mean don't run MultiChannel.
- OperChnCap: The current value for the Number of Channels supported parameter. This value is included as the number of channels supported parameter in the local MCh TLV. The range for this variable is 1-0xffe.
- OperSVIDList: The current value for (channel,SVID) assignments. This is the list of (Channel,SVID) pairs included as the (Channel,SVID) pairs in the local MCh TLV. The total size of the list may not exceed 167 pairs. If the list is empty its value is NONE. The valid range for each channel of this list is from 1-0xffe. The valid range for each SVID in the list is from 1 to 0xfff. When the SVID is value is 0xfff the SVID is unconfigured. For the S role a SVID of 0xfff indicates a request for a channel. For the B role an SVID of 0xfff indicates an unconfigured channel.

Remote Operational Variables

- RemoteVersion: The current value for the remote MCh version parameter. This value is included as the Version parameter in the remote MCh TLV. NULL means no remote MCh TLV exists in the local LLDP database. The value for this variable may be VER0=100b setting any other value will result in stopping MultiChannel operation.
- RemoteChnCap: The current value for the Number of Channels supported parameter. This value is included as the number of channels supported parameter in the remote MCh TLV. NULL means no remote MCh TLV exists in the local LLDP database. The range for this variable is 1-0xffe.
- RemoteSVIDList: The current value for (channel,SVID) assignments. This is the list of (Channel,SVID) pairs included as the (Channel,SVID) pairs in the remote MCh TLV. NULL means no remote MCh TLV exists in the local LLDP database. If the list is empty but the MCh TLV is present its value is NONE. The total size of the list may not exceed 167 pairs. The valid range for each channel of this list is from 1-0xffe. The valid range for each SVID in the list is from 1 to 0xfff. When the SVID is value is 0xfff the SVID is unconfigured. For the S role a SVID of 0xfff indicates a request for a channel. For the B role an SVID of 0xfff indicates an unconfigured channel.
- LastRemoteSVIDList: Temporary local copy of the RemoteSVIDList. This variable is not included in the MCh TLV. The LastRemoetSVIDList has the same syntax as RemoteSVIDList.

MultiChannel State Machine

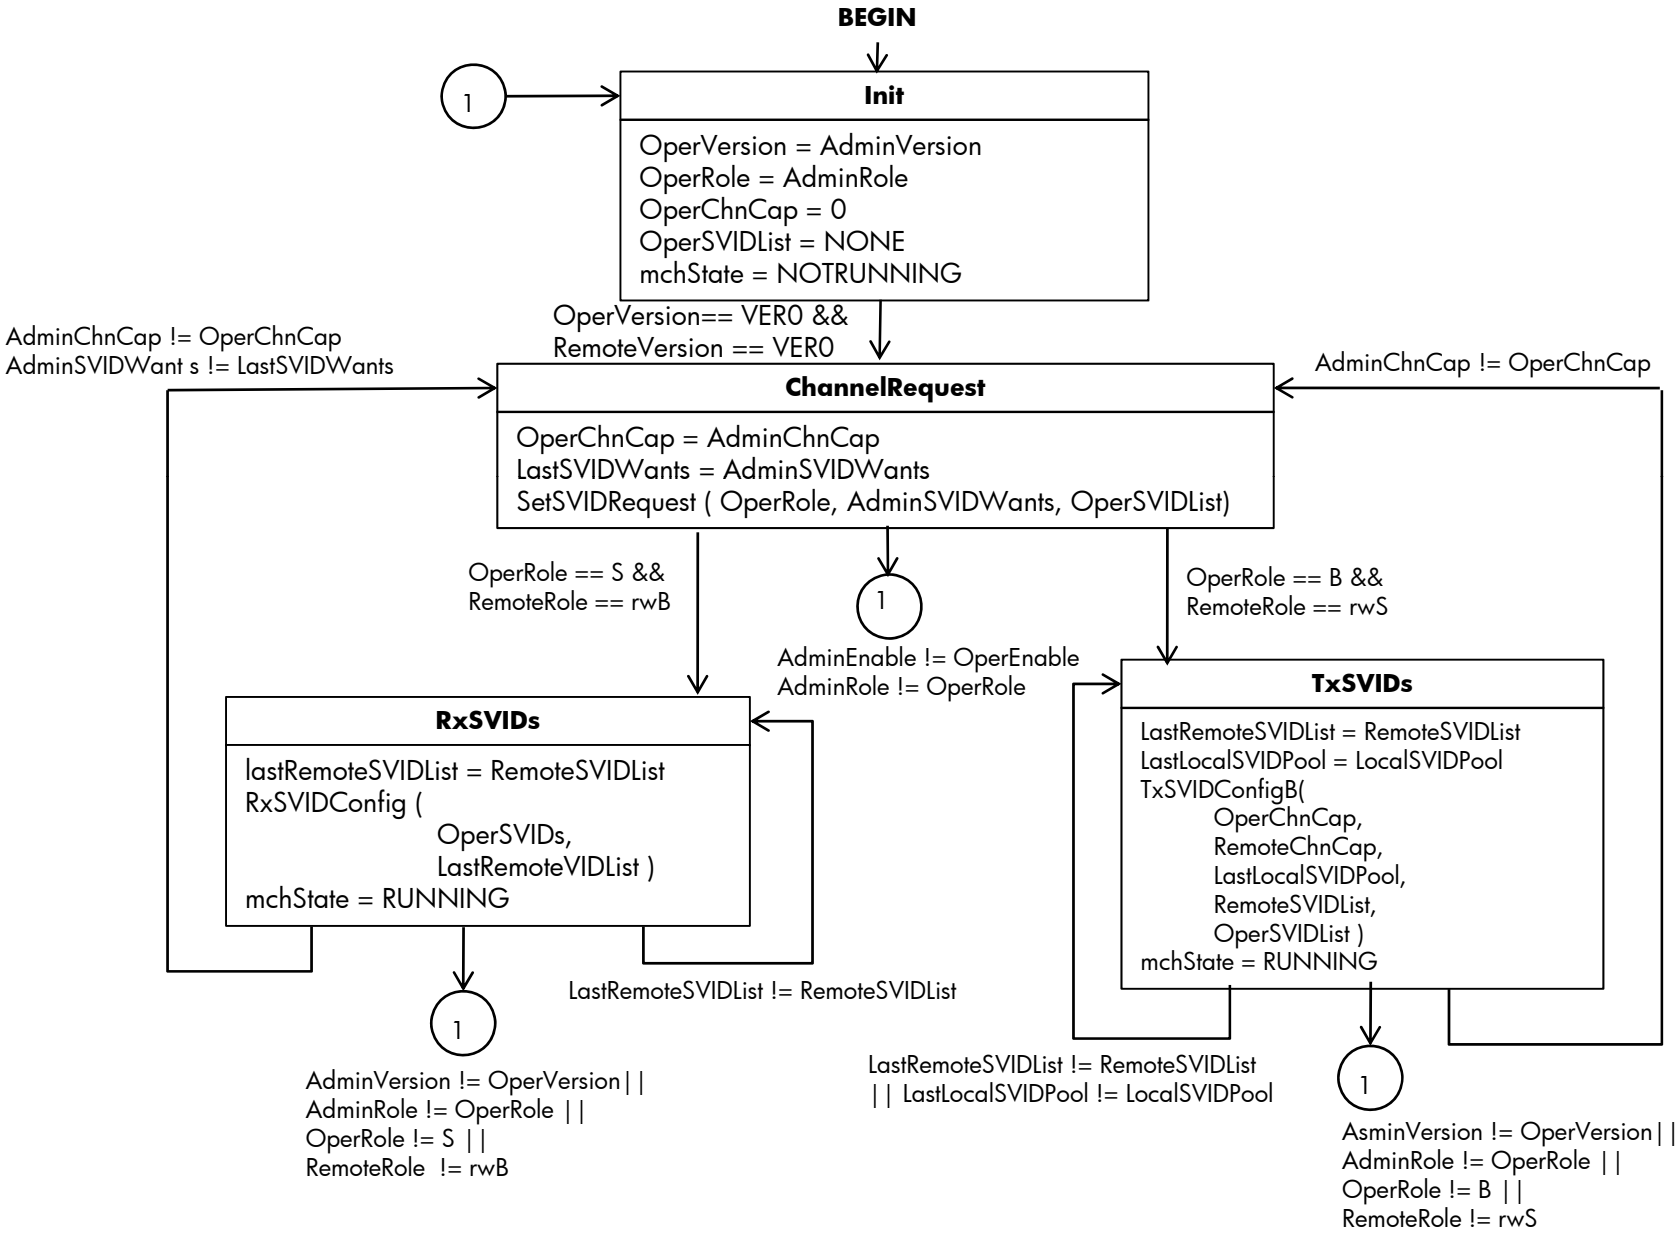

State Machine Functions

- SetSVIDRequest(OperRole, AdminSVIDWants, OperSVIDList)
  - This function creates the OperSVIDList placed in the Local TLV database.
  - If the OperRole for the equipment is R then the OperSVIDList remains unchanged.
  - If the OperRole for the equipement is S two possible cases exist. In the first case we don't have any configured channels, indicated by OperSVIDList being equal to NONE. In this case the function places the AdminSVIDWants in OperSVIDList. In the second case we already have a running configuration indicated by the OperSVIDList not equal to NONE. In this case the function compares the AdminSVIDWants with the OperSVIDList. All active channels in the OperSVIDList which are in the AdminSVIDWants are kept active and in addition any channels not currently in the OperSVIDList are requested by including them in the OperSVIDList along with a 0xfff SVID number.
- RxSVIDConfig (OperSVIDs, LastRemoteVIDList)
  - This function creates the OperSVIDList placed in the Local TLV database for an S role equipment
  - The function compares the AdminSVIDWants with the LastRemoteSVIDList. For each AdminSVIDWants channel with an SVID assignment in the LastRemoteSVIDList a (Channel,SVID) pair is generated in the OperSVIDList. For each AdminSVIDWants channel without an SVID assignment in the LastRemoteSVIDList a (Channel,0xfff) pair is generated in the OperSVIDList.
- TxSVIDConfigB(OperChnCap, RemoteChnCap, LastLocalSVIDPool, RemoteSVIDList, OperSVIDList)
  - This function creates the OperSVIDList placed in the Local TLV database for an S role equipment
  - First the function takes the smaller of the OperChnCap and RemoteChnCap and truncates the RemoteSVIDList to the smaller of the two.
  - A new OperSVIDList is created as follows:
    - For each channel in the RemoteSVIDList with a (channel,SVID) pair in the OperSVIDList the (channel,SVID) remains unchanged unless the SVID is no longer part of the LastLocalSVIDPool. If the SVID is no long in the pool a new one is selected if available. If no SVID is available the (channel,SVID) pair will be deleted from the OperSVIDList.
    - For each channel in the RemoteSVIDList without a (channel,SVID) pair in the OperSVIDList an SVID is obtained from the LastLocalSVIDPool (the pool for Bridge resources) if available. If no SVID is available the (channel,SVID) pair will be deleted from the OperSVIDList.

Summary

- MDP allows configuration of S-VLANs used to simplify the operation of the station.
- MDP is a LAN level LLDP TLV exchange protocol used to discover and configure the MultiChannel capability
- MDP operates by assigning the Station and Bridge an S or B role. For a successful MultiChannel configuration one side must be S and the other B.
- During configuration the S side requests channel resources and the B side supplies channel resources.

BACKUP SLIDES

Static Configuration Success (new role)

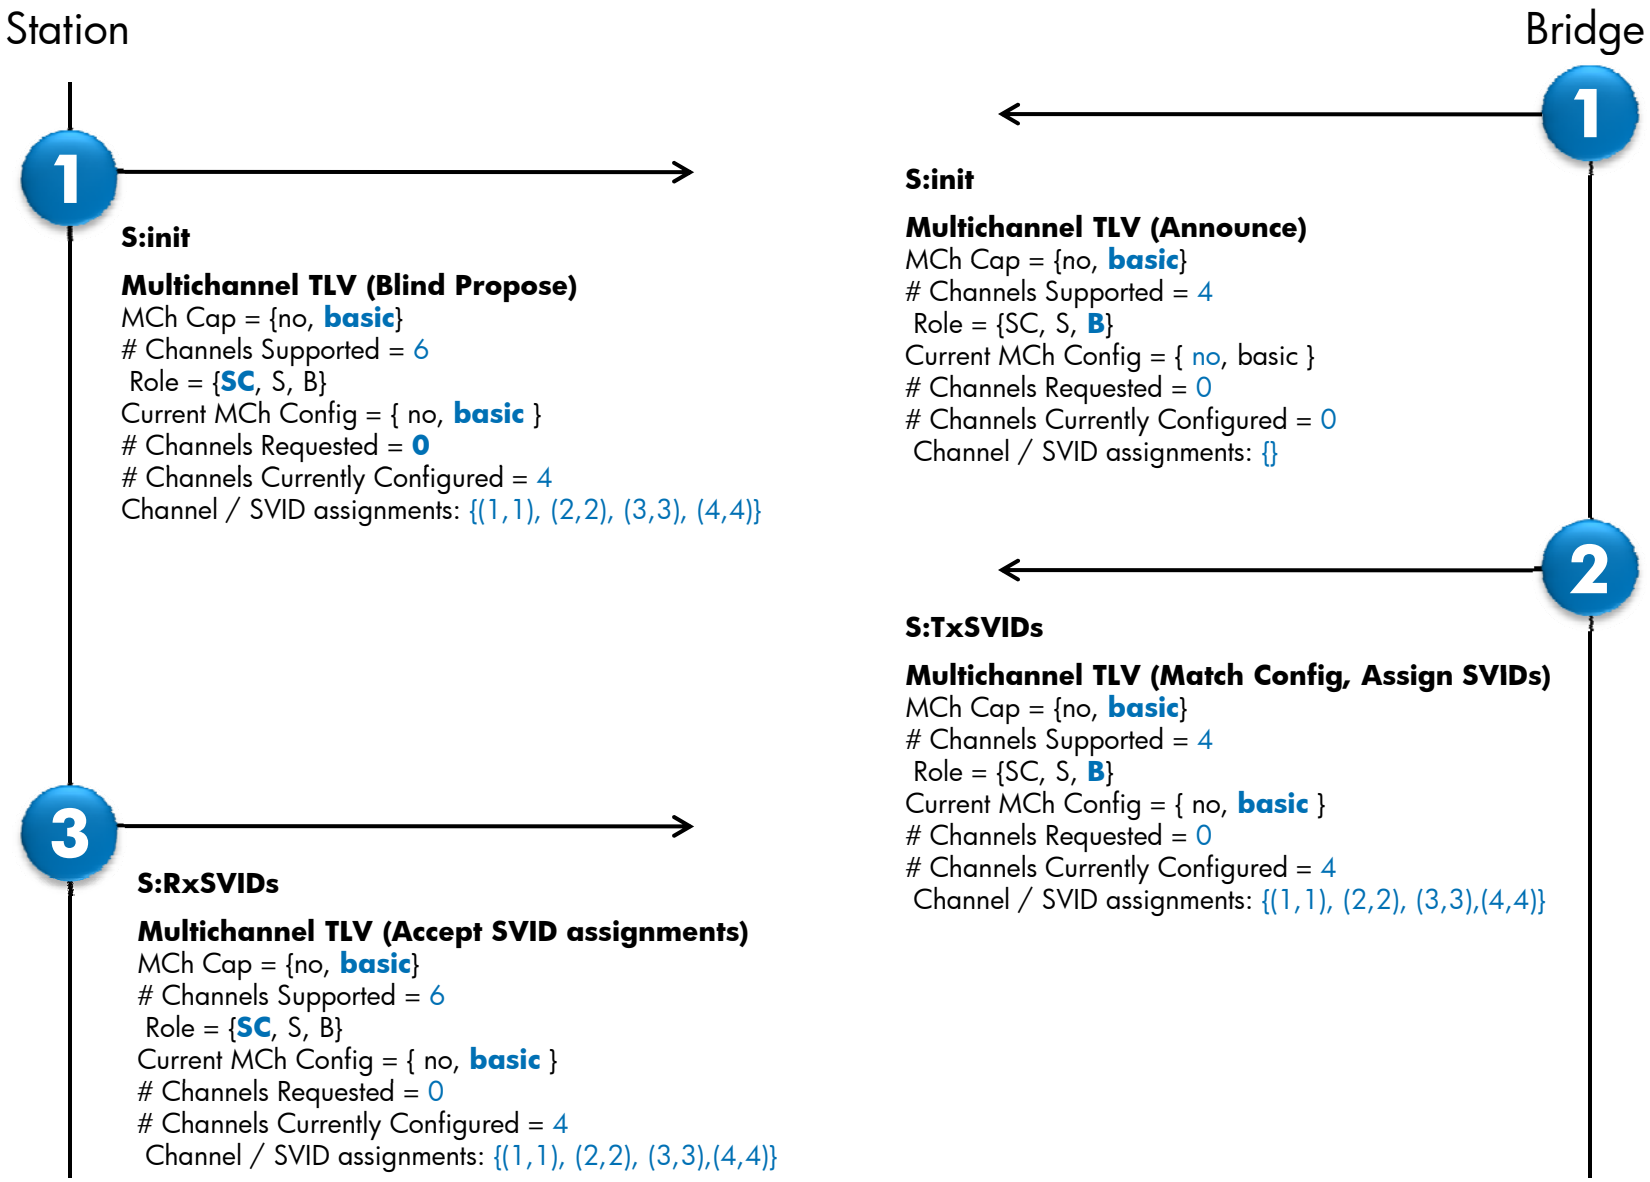

Static Configuration Fail (new role)

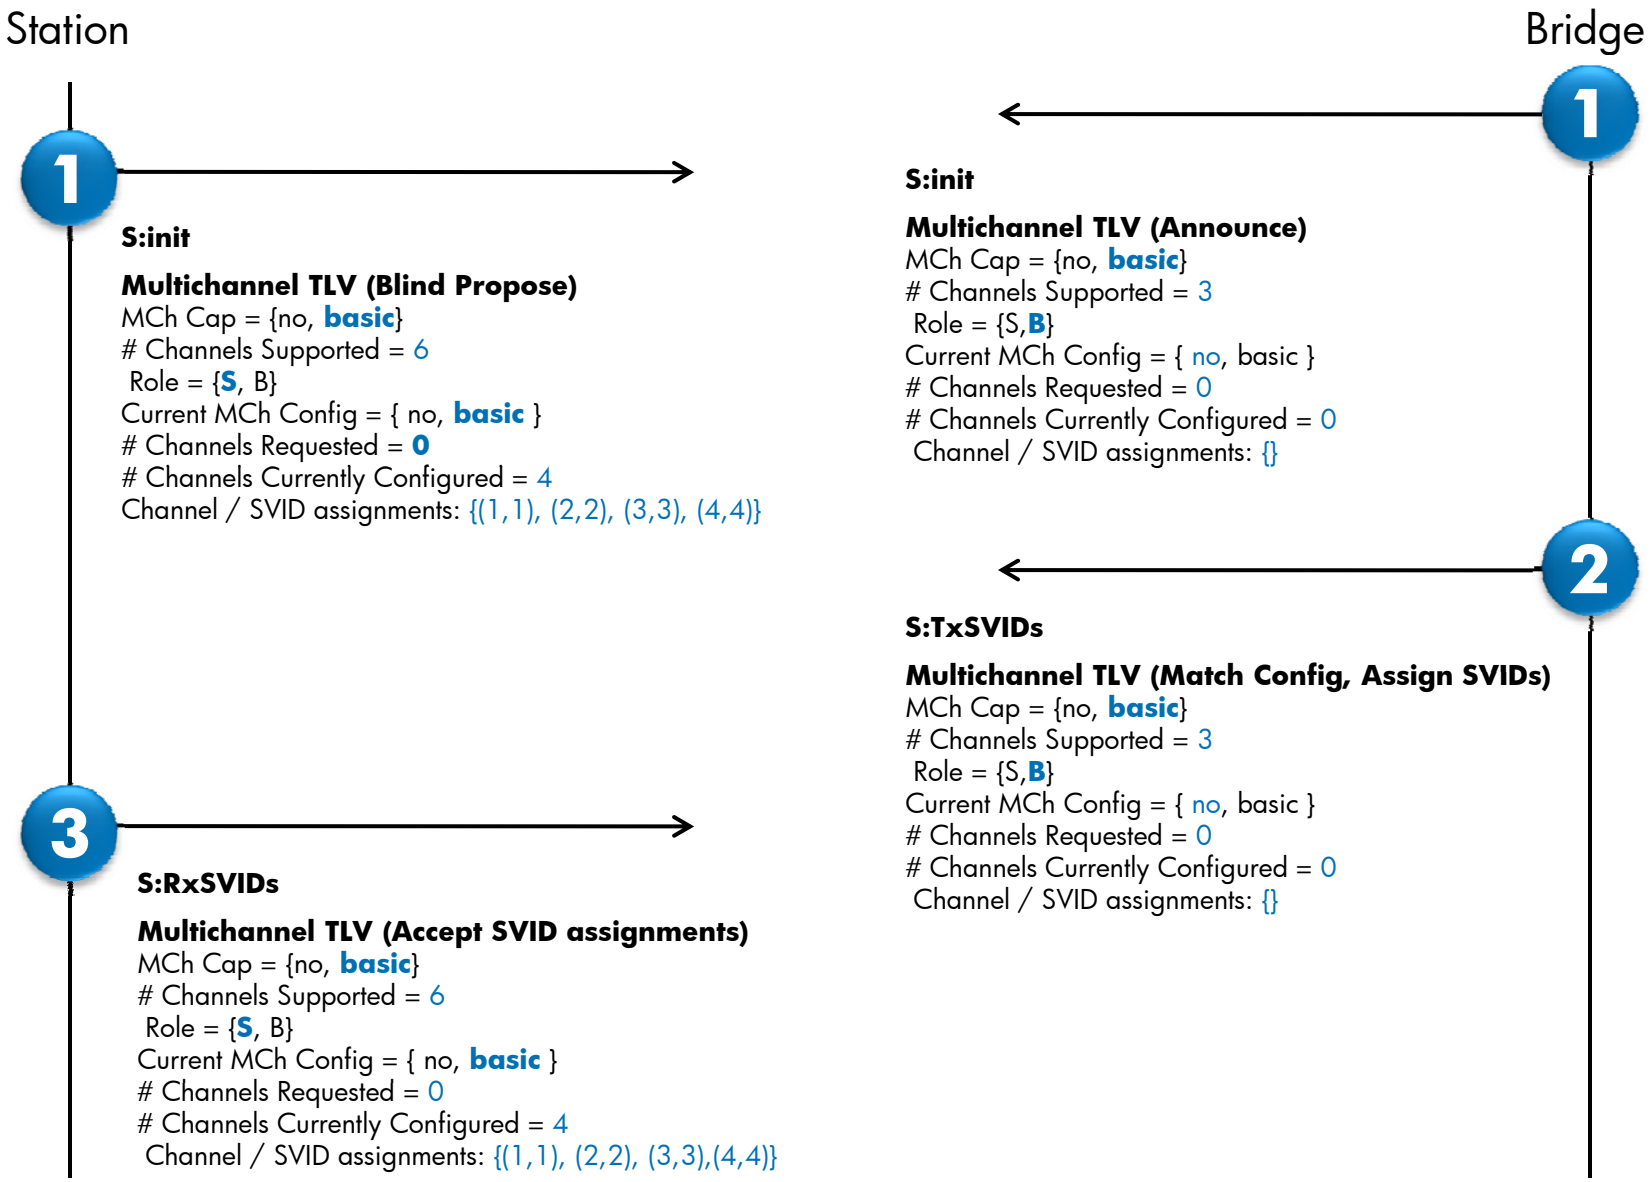

Startup-No MCh —X→ Startup-MCh

- LAN LLDP
  - EVB
- ETTP
  - VDP

- LAN LLDP
  - MDP
- Per Channel
  - LLDP
    - EVB
  - ETTP
    - VDP

- Transition between No MCh and MCh may be disruptive.