FUJITSU

The Emerging Optical Control Plane

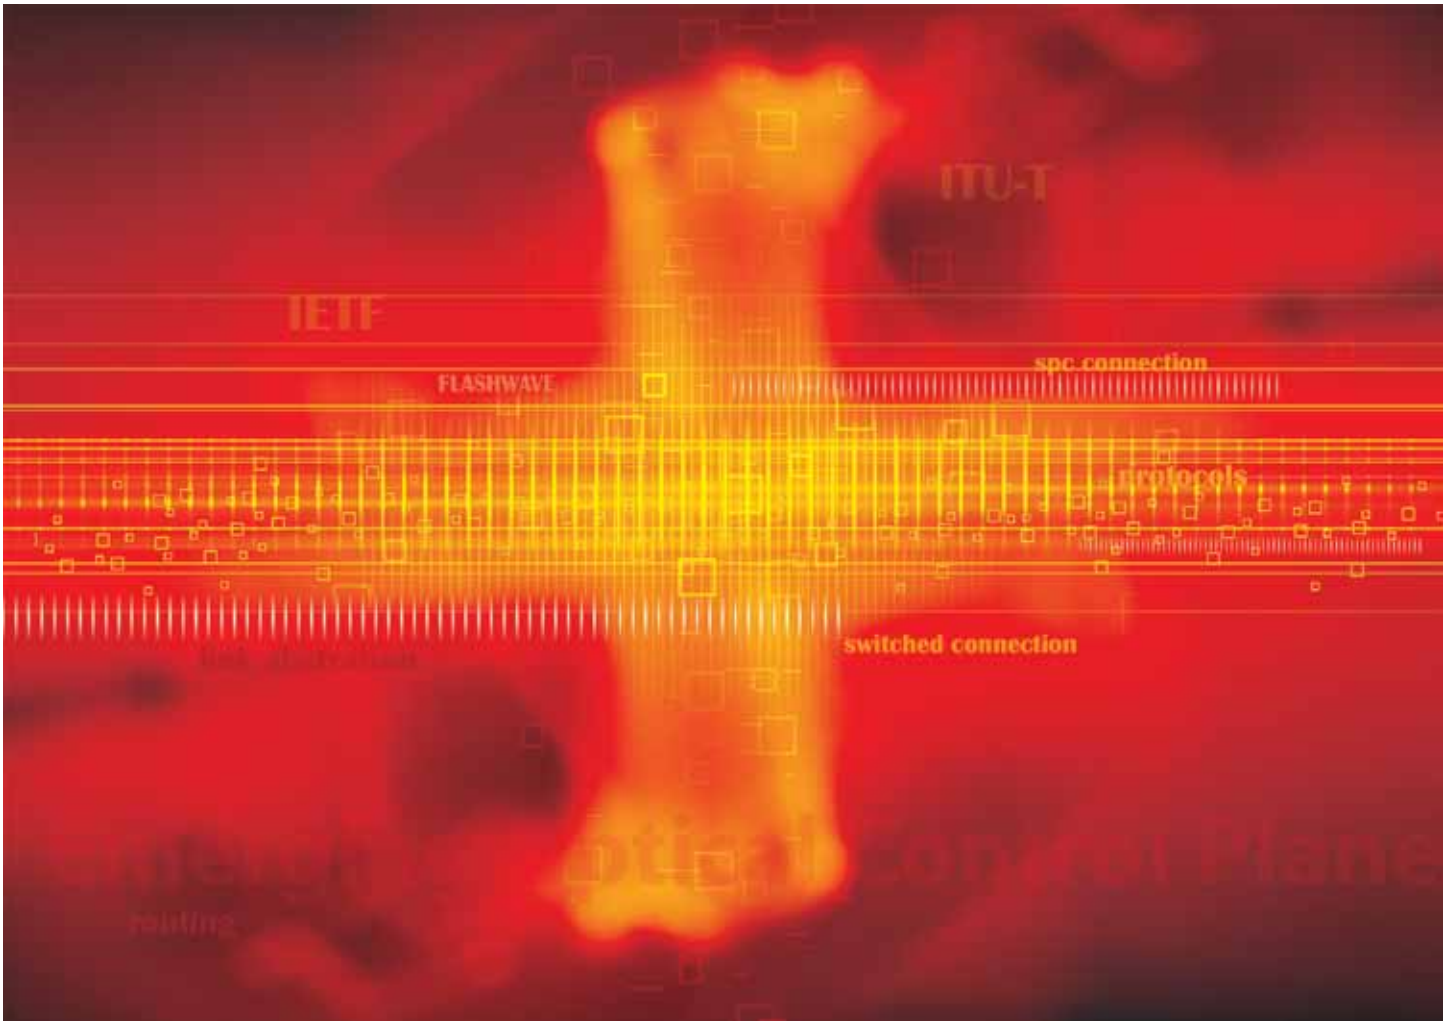

Traditional transport networks can be modeled as the interaction of two operating planes: a transport plane and a management plane. In this model, the transport plane carries the user data and comprises network equipment, such as line interface cards, switch fabrics, backplanes and fiber plant. Network OAM&P is fully handled by the management plane, implemented by an EMS, NMS, and/or OSS.

Now, we are beginning to see the deployment of optical control planes that sit between the management and transport planes (Figure 1). The control plane moves some of the network intelligence down to the NEs. As a result, the NEs have access to complete network topology and resource information, and can use this to plan, establish and maintain user services.

Figure 1: Operating Planes

In the optical transport field, the control plane has been called different things by different standards bodies. In the IETF, the optical control plane is referred to as Generalized Multi-Protocol Label Switching, or GMPLS [1]. Within the ITU-T, the optical control plane initiative is referred to as Automatic Switched Optical Network (ASON) [2]. In general, the ITU-T has defined the architecture and requirements for an optical control plane, whereas the IETF has mostly focused on developing the control plane protocols. A third standards body, the OIF, has been working on applying the IETF GMPLS protocols to the ITU-T ASON architecture, in order to promote multivendor interoperability [8, 13, 14, 15, 16, 17, 18]. The standards bodies applicable to the optical control plane are shown in Figure 2.

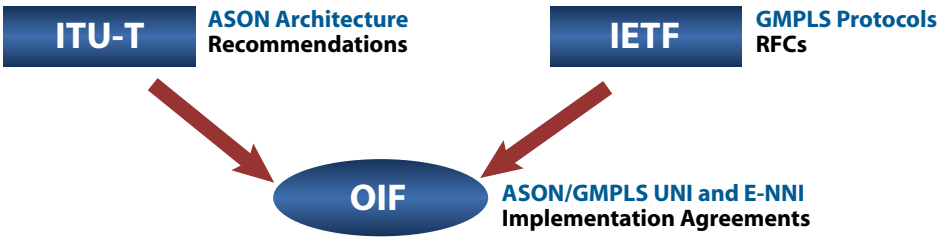

Figure 2: Standards Bodies Developing for the Optical Control Plane

The control plane offers a number of benefits to carriers. A control plane-enabled network can support new, dynamic services. These services include bandwidth-on-demand applications, customer-initiated service requests via a UNI, and scheduled or time-of-day based services. The control plane also integrates well with mesh networks. Coupled with dynamic service restoration, the control plane can improve network efficiency and resilience. Finally, the control plane can help reduce operating expenses by decreasing operator workload, reducing service turn-up fallout, and supporting multivendor and multicarrier interoperability.

This paper begins by discussing some general control plane concepts, such as call and connection control, network reference points and connection types. The control plane applications (discovery, routing, path computation, signaling) are covered, along with the protocols used in the FLASHWAVE® products to support these applications. The paper then addresses the SCN and concludes with a discussion of network applications.

Control Plane Concepts

This section describes important control plane concepts, including call and connection control, in addition to the ITU-T network reference points and connection types.

Call and Connection Control

ITU-T architecture supports the concept of calls and connections (Figure 3). A call is an association between the end users to support a service. A call controller is used to signal information between the user and the network to control the setup, modification, and teardown of the connections that establish the service. Call controllers may also be required at domain boundaries to provide support for call/connection control within the network domains.

A call is broken down into a series of call segments, each of which defines the association between two adjacent call controllers. Typically, the network domain boundaries define the scope of a call segment. The call segment is responsible for controlling the connections within its domain. The call controller for the segment determines the type and number of connections required to meet each service request. For instance, to meet a service request's availability requirement, a call segment may establish a 1+1 protected service by creating two connections: a working connection and a protect connection. Another call segment may establish a single connection and rely on dynamic restoration for protection. The decision on how to instantiate a service is determined by local provisioning (policy).

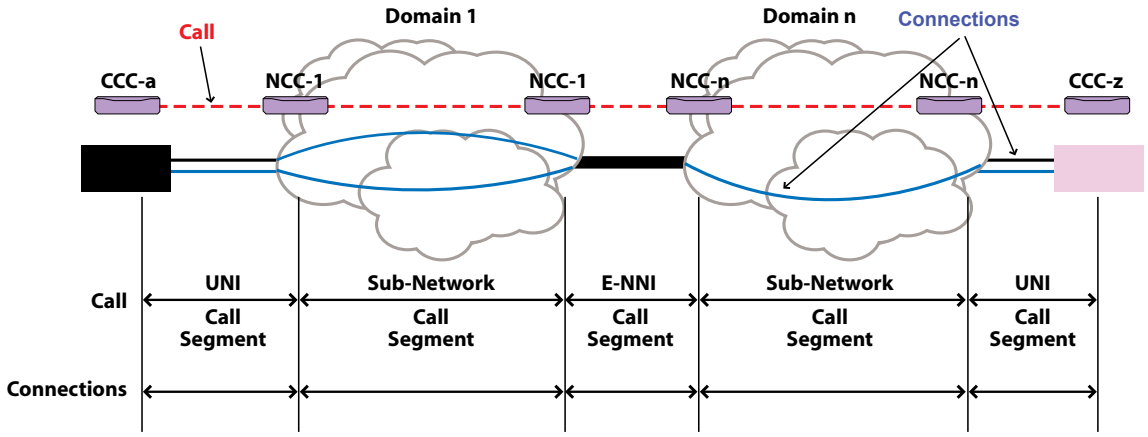

Figure 3: Call and Connections (from OIF-E-NNI Sig-01.0) [13]

Network Reference Points

ITU-T defines three reference points within the network: the User-Network Interface (UNI), the Internal Network-Network Interface (I-NNI), and the External Network-Network Interface (E-NNI). These reference points are shown in Figure 4.

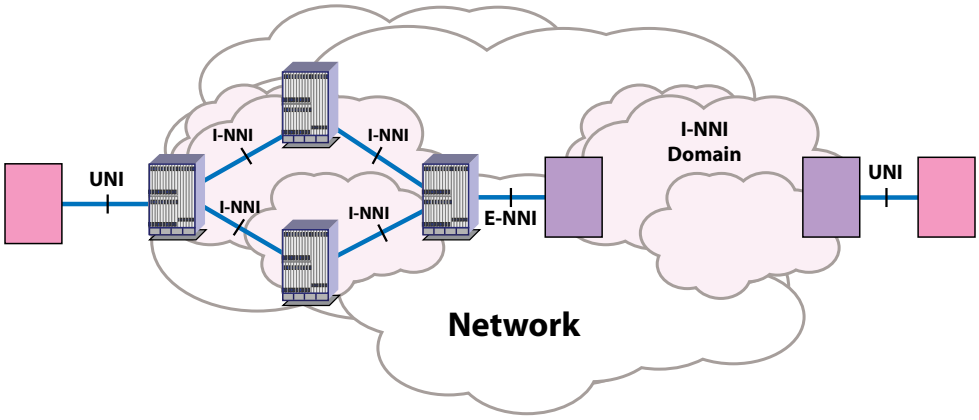

Figure 4: ITU-T Network Reference Points

The UNI reference point defines the boundary between a client and the network. The client requests network services over the UNI interface, and the network responds by establishing calls and connections to meet the request.

The NNI defines the interface between devices within the network. The NNI can be further divided into either an internal or external interface. The internal NNI, or I-NNI, is typically a single-vendor interface contained within a single-carrier network. Since it is single-vendor, this interface may contain proprietary elements specific to that vendor. It also assumes a full trust model and so maximizes the information exchange between devices.

The external NNI, or E-NNI, is an external interface between devices within the network that crosses domain boundaries. Domain boundaries are defined by the carriers and can include administrative boundaries within a carrier's network, boundaries between different vendors within a carrier's network, or boundaries between carriers. The information exchanged across the E-NNI is usually more restricted than that exchanged across the I-NNI. For example, topology information may be abstracted for scalability or to hide the internal details of a carrier's network. Whereas the I-NNI may have the proprietary elements, the E-NNI is standardized to allow for multivendor interoperability.

As mentioned previously, the OIF has been developing interoperability standards for the optical control plane. More specifically, the OIF has been developing standards for the UNI and E-NNI reference points. The OIF does not attempt to standardize the I-NNI reference point, as this is typically a single-vendor domain.

Types of Connection

The ITU-T defines three connection types. These connection types distinguish the distribution of the connection management functions between the management and control planes:

- **Permanent Connections (PC)** – The management system establishes the connection directly using cross-connect provisioning on each device within the network. In this case, there is no control plane involvement.
- **Switched Connection (SC)** – The customer equipment (a router, for example) triggers establishment of the connection using control plane signaling over the UNI interface. The end-to-end connection is established using the control plane without management plane involvement. The SC connection type is shown in Figure 5.

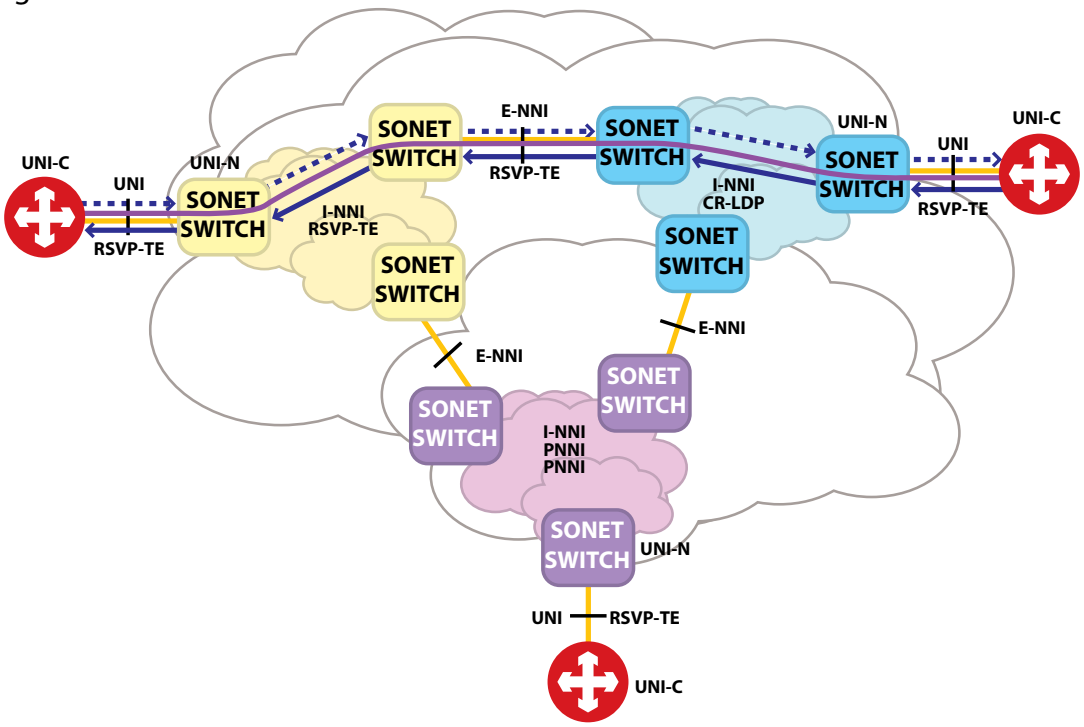

Figure 5: SC Connections

- **Soft Permanent Connection (SPC)** – The SPC connection type is a hybrid that involves both the management and control planes to establish the connection. The management system establishes the connection between the client equipment and the network as a permanent connection. The management system also triggers the establishment of the connection within the network by making a call request to the ingress node. The ingress node uses NNI control plane signaling to establish the connection within the network. The SPC connection type is shown in Figure 6.

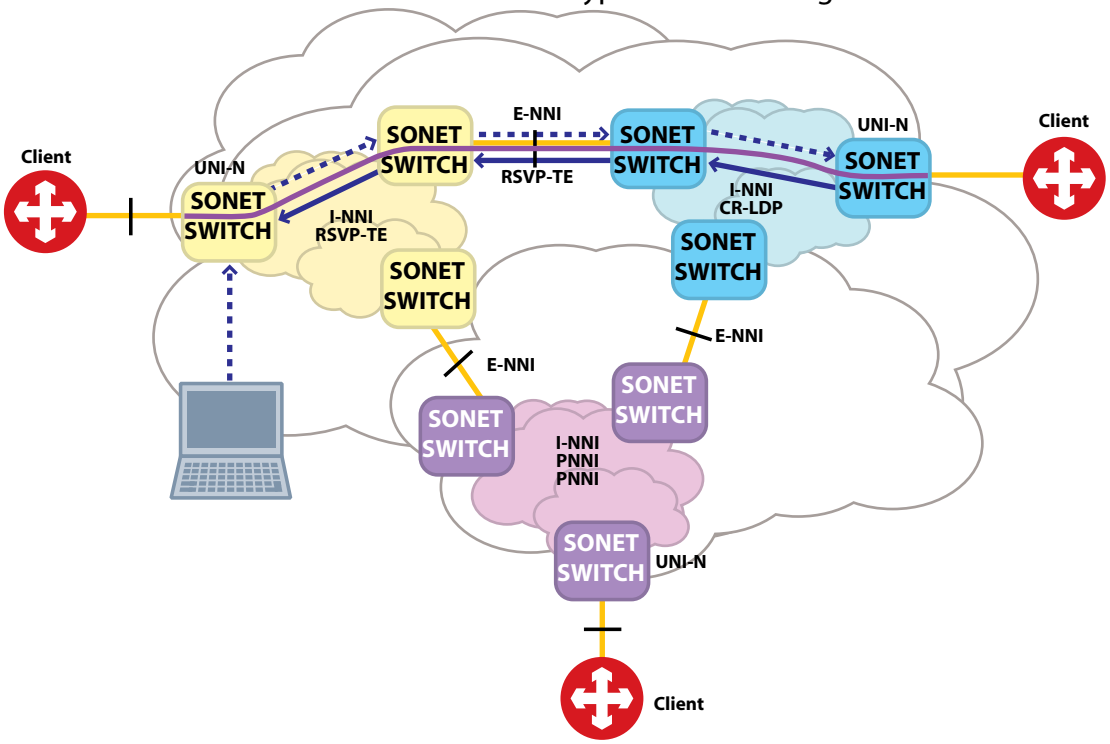

Figure 6: SPC Connections

Control Plane Applications and Protocols

The control plane can be seen as a set of applications that support the ability to establish a service through the network. This is in contrast to the traditional method, where a management system establishes the service by provisioning cross-connects on each individual network element.
The control plane applications are:

- Discovery
- Routing
- Path Computation
- Signaling

Discovery Application

The discovery application is responsible for the discovery of neighbors and the links between neighbors. Neighbor discovery is automated when there is an in-fiber IP communication channel (i.e., DCC, GCC, OSC, etc.) to that neighbor. The local node sends a discovery message over the in-fiber channels. Neighbors receiving this message respond to the originating node, completing the neighbor discovery process. Node-level information is exchanged during this process to allow for the discovery of the neighbor's Node IP address and node identifier.

The discovery protocol also supports the discovery of link connectivity between neighbors. This is accomplished through the link verification process. Link verification works by sending a unique test message over the link being discovered. The local node sends a discovery request to its neighbors to search for that discovery message. The neighbor node that sees this message on an incoming link responds to the originating node. In this manner the local and remote link identifiers are discovered. The local node repeats this process for each of its links.

The discovery test message is technology-specific and uses either the overhead or payload bytes of the in-fiber signal. For SONET links, link verification uses the J0 section trace. For WDM links, the test message is a proprietary format sent over the OSC.

Link connectivity verification also verifies that the transmit and receive ports on the two nodes are consistently fibered. If not, a fiber mismatch alarm is transmitted to the user.

The discovery application maintains neighbor and link connectivity information and updates this in real time. This information is used by the routing application to distribute topology information throughout the network.

Routing Application

The routing application is responsible for propagating the link connectivity information to all nodes within the network. This results in the formation of the TE database. The TE database contains the information necessary to determine the network topology, as well as resource information to support traffic engineering (e.g., link bandwidth availability).

As the number of nodes increases, scalability of the routing protocol eventually becomes an issue. Hierarchical routing addresses this issue. A single-level hierarchical routing model is defined in OIF E-NNI routing 1.0, as shown in Figure 7. In this model, a separate and independent instance of a routing protocol is run within each I-NNI domain (considered Level 0). These I-NNI domains may run different routing protocols, perhaps even vendor proprietary implementations. For instance, in Figure 7, Domain 1 and Domain 3 run OSPF-TE, while Domain 2 runs ISIS-TE.

Routing information (i.e., topology, resource, and reachability information) is exported from the Level 0 I-NNI domains into the E-NNI routing area (considered Level 1). At the Level 1 E-NNI domain, all routing controllers must run the same routing protocol for interoperability. OIF E-NNI Routing 1.0 mandates that the Level 1 E-NNI routing protocol be OSPF-TE.

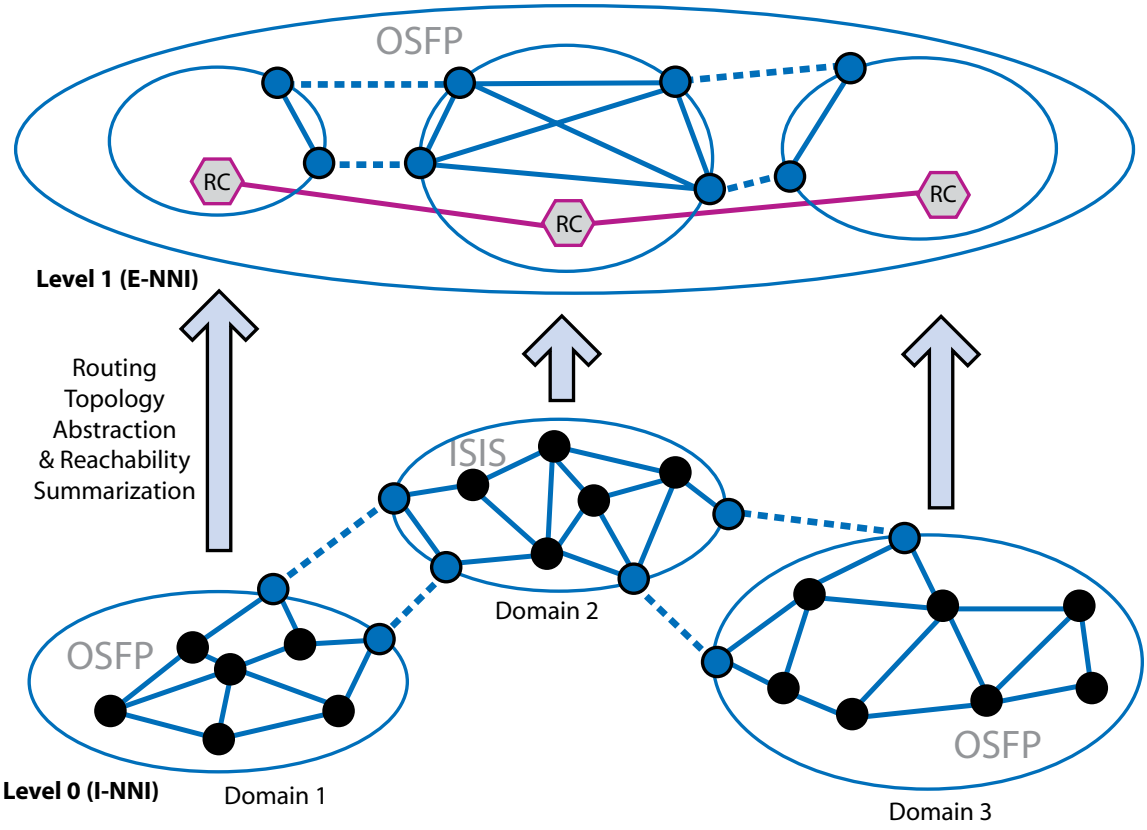

Figure 7: Single-Level Hierarchical Routing Model

The routing hierarchy architecture allows for additional levels of hierarchy (Level 2 and above). Although the current OIF routing standard only defines a single hierarchical level (Level 1), it is expected that future standards development will extend the routing protocol to support additional hierarchical levels.

During the export of routing information from the Level 0 I-NNI domain into the Level 1 E-NNI domain, topology information can be abstracted, and reachability information can be summarized. This abstraction/summarization can drastically reduce the amount of routing information carried within the E-NNI routing domain, thus improving network scalability. Some different abstraction methods are listed below and shown in Figure 8:

- **No Abstraction** – Topology information is imported into the upper hierarchical level without change.
- **Link Abstraction** – Domain border nodes are imported into the upper hierarchical level without change. However, instead of showing the actual internal topology, the domain is represented as virtual links connecting the domain border nodes. The virtual topology can be a full or partial mesh between border nodes.
- **Pseudo-Node Abstraction** – Similar to link abstraction as the domain border nodes are imported into the upper hierarchical level. But instead of showing link connectivity between border nodes, the internal topology is represented as a connection from the border nodes to a logical pseudo-node that represents the internal network.
- **Node Abstraction** – The entire level 0 domain is abstracted as a single logical node.

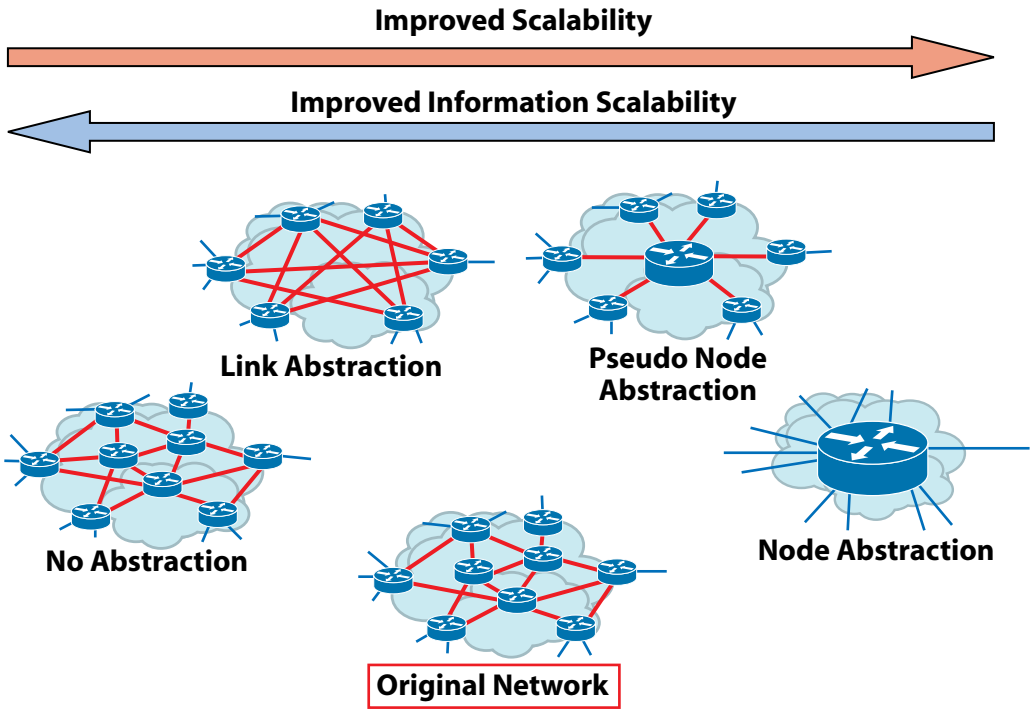

Figure 8: Routing Abstraction Methods

The various abstraction mechanisms allow for a trade-off between information loss and scalability. Increased information loss helps to achieve scalability, but it can make certain functions more difficult, such as diversity or end-to-end latency calculation. On one end of the spectrum, "no abstraction" provides for complete topology information into the level 1 routing instance, but confers no scalability advantage. On the other end, "node abstraction" completely loses the internal topology information but provides the best scalability.

Path Computation Application

The path computation application performs the planning function for the control plane. Path computation may be requested to determine one or more routes through the network in response to a service request. Path computation provides the necessary service information, such as the source ingress location, destination egress location and service constraints. The TE database is used to compute routes during path computation.

Each link in the TE database has a cost attribute. The cost of a route is the sum of the cost of each link that makes up the route. A CSPF algorithm is used to find the lowest cost route from source to destination that meets the service constraints. Examples of service constraints that can be considered by path computation are bandwidth requirements, lambda continuity for WDM circuits, maximum allowable end-to-end latency, and include/exclude links or nodes.

The path computation function can also calculate diverse routes for dedicated protected services. Path computation supports link, node, and SRLG diversity.

Signaling Application

The signaling application is responsible for setting up, modifying, and tearing down end-to-end services. When the ingress node receives a service request, the signaling controller requests an optimal path from the path computation application that meets the service constraints. The signaling controller then proceeds to establish the service using the signaling protocol.

The signaling controller associates a call identifier to identify the call associated with the service. This call identifier has end-to-end scope and does not change as it crosses domain boundaries. All connections associated with the call share the same call identifier.

For each call segment, the ingress signaling controller for that call segment associates a connection identifier for the connection that exists within the call segment. Each call segment assigns the connection identifiers independently and this identifier is only valid within the call segment to which it pertains. Thus, for an end-to-end service, the connection identifier may change as the service passes domain boundaries (such as across an E-NNI interface).

The signaling protocol supports alarm-free connection setup by using a multipass setup process, as shown in Figure 9. The call ingress node starts by sending a connection setup request to the egress node. This request is used by intermediate nodes to verify admission criteria, check bandwidth availability, and reserve resources. The egress node, upon accepting the connection request, sends a connection request indication back to the ingress node. This indication triggers nodes to establish the cross-connects in their switch matrix, but in an alarm-suppressed state. Finally, the ingress node finalizes the setup by sending a confirmation message that transitions the connection to the alarm-enabled state.

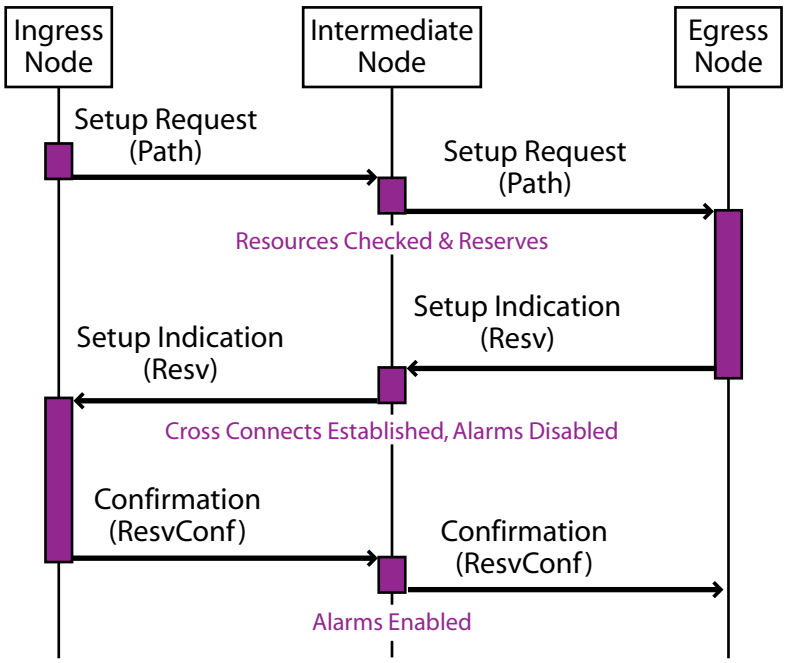

Figure 9: Connection Setup

Connection teardown is also alarm-free. A two-pass connection teardown is used as shown in Figure 10. To initiate the teardown, the call ingress node sends a connection release request. This message triggers nodes to suppress alarms for the connections. The egress node responds by sending a connection release indication. Upon receiving the release indication, nodes release the cross-connects associated with the connection.

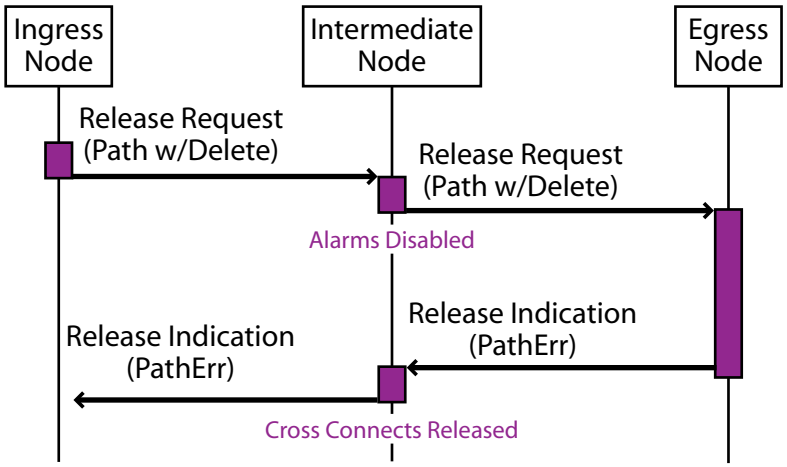

Figure 10: Connection Teardown

Application Interaction

The control plane applications work together to automatically establish services within the network. This interaction is shown in Figure 11. The discovery application discovers the local links connecting to neighbors. This information can also be manually provisioned if discovery is not enabled or supported. The routing application broadcasts this information to all nodes in forming the TE Database. When a user requests a service from the network, path computation is called to compute a route using the TE database. This route is provided to signaling to establish the service. A user can also explicitly provide the route for the service, in which case signaling is immediately initiated without path computation.

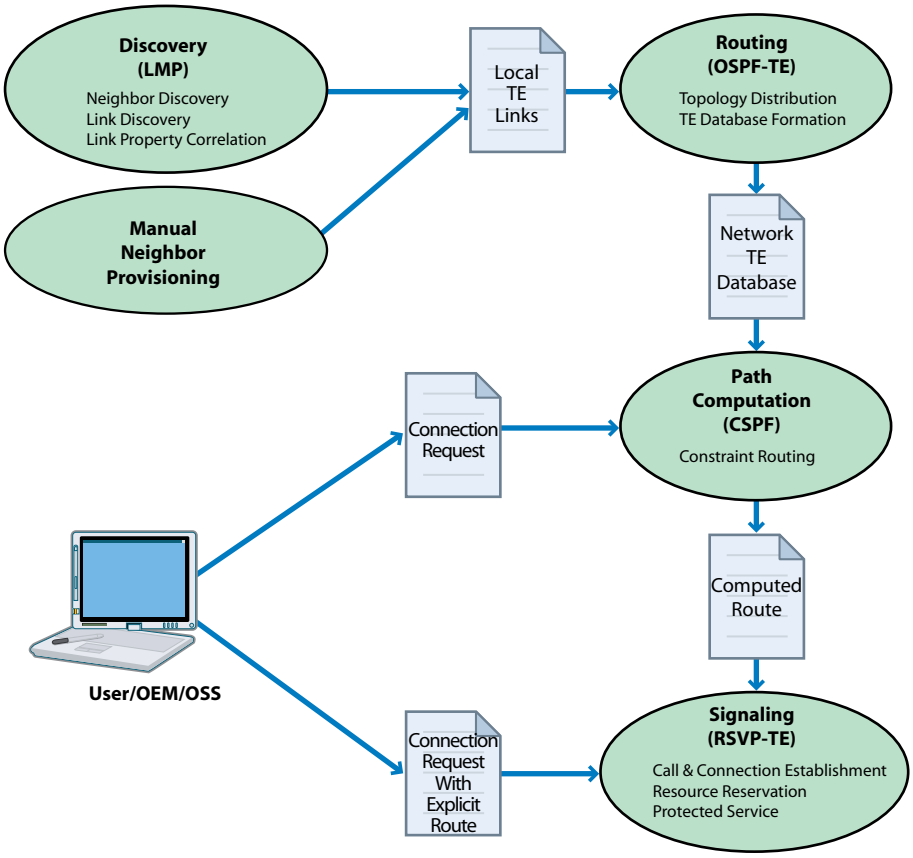

Figure 11: Control Plane Application Interaction

The Fujitsu Approach to Control Plane Protocols

The Fujitsu control plane approach is shown in Figure 12.

- **Discovery** – LMP is used as the discovery protocol [3]. Discovery is supported over the I-NNI, E-NNI and UNI reference points.
- **Routing** – OSPF-TE is used as the routing protocol at the I-NNI and E-NNI level [4, 5, 6, 7]. At the E-NNI level, OIF routing extensions are supported for interoperability [8]. Routing is supported over the I-NNI and E-NNI reference points. The UNI reference point does not support routing because carriers do not want to export their network topology information outside their domain.
- **Signaling** – RSVP-TE is used as the signaling protocol [9, 10, 11, 12]. OIF E-NNI 2.0 extensions are supported for interoperability [18]. Signaling is supported over the I-NNI, E-NNI and UNI interfaces.

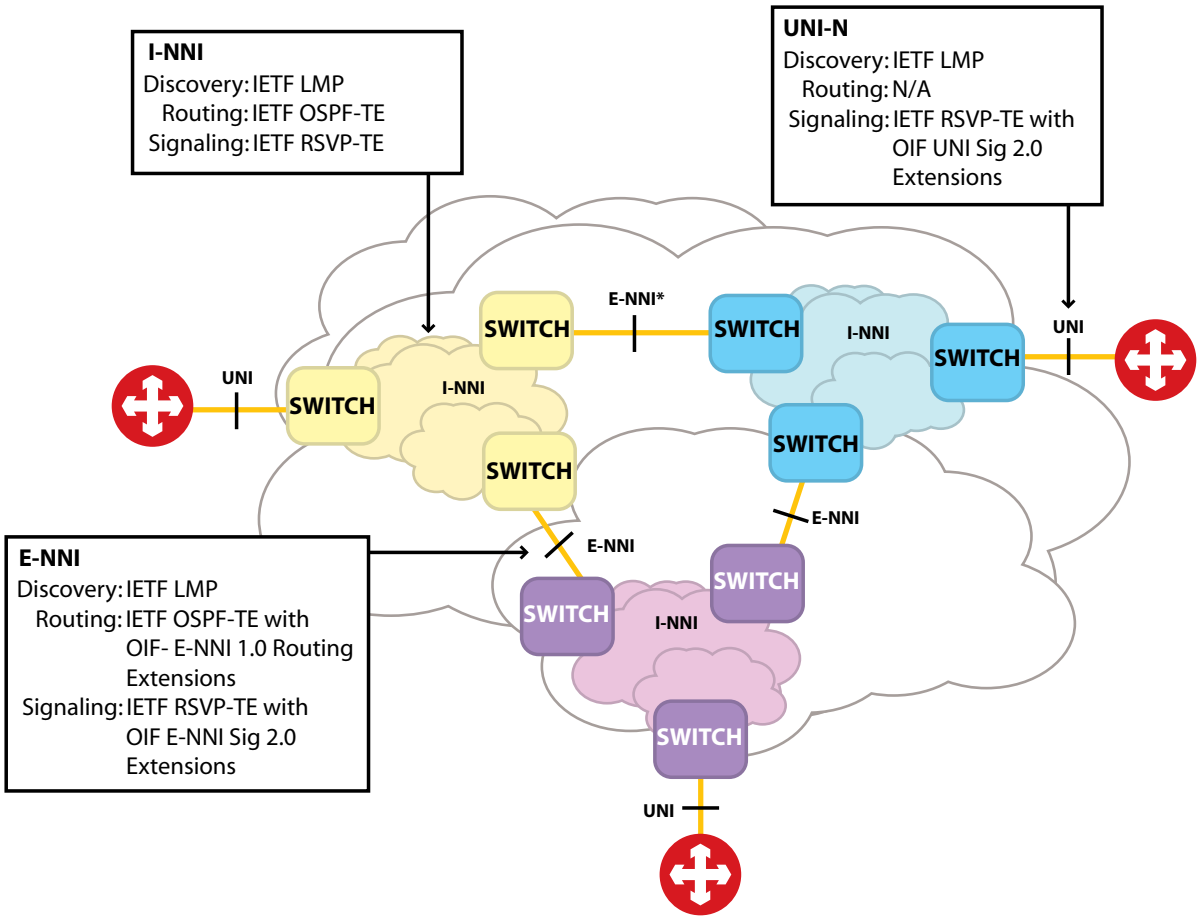

Figure 12: The Fujitsu Control Plane Approach

Signaling Communication Network

The DCN can be broken down into an MCN and an SCN. The MCN is the traditional part of the DCN, and provides the communications infrastructure for management systems to communicate with the network elements. The SCN provides the communication infrastructure for the control plane. As the control plane protocols are IP-based, the SCN requires an IP routing protocol. Fujitsu FLASHWAVE products support Integrated IS-IS as the SCN IP routing protocol. As an additional option, some products may also support OSPF.

SCN Topology

The SCN topology can be congruent to the transport network by using embedded communication channels (e.g., DCC, GCC or OSC). The SCN topology can also be non-congruent by using out-of-band facilities (e.g., LCN port) in addition to (or in place of) the embedded channels. Figure 13 shows a sample network layered by the SCN and control plane topology database. In the SCN layer, nodes n1, n2, n3 and n4 are interconnected via DCC as are nodes n5 and n6 and nodes n7, n8 and n9. Nodes n3, n5 and n7 are connected to a LAN via their respective LCN ports.

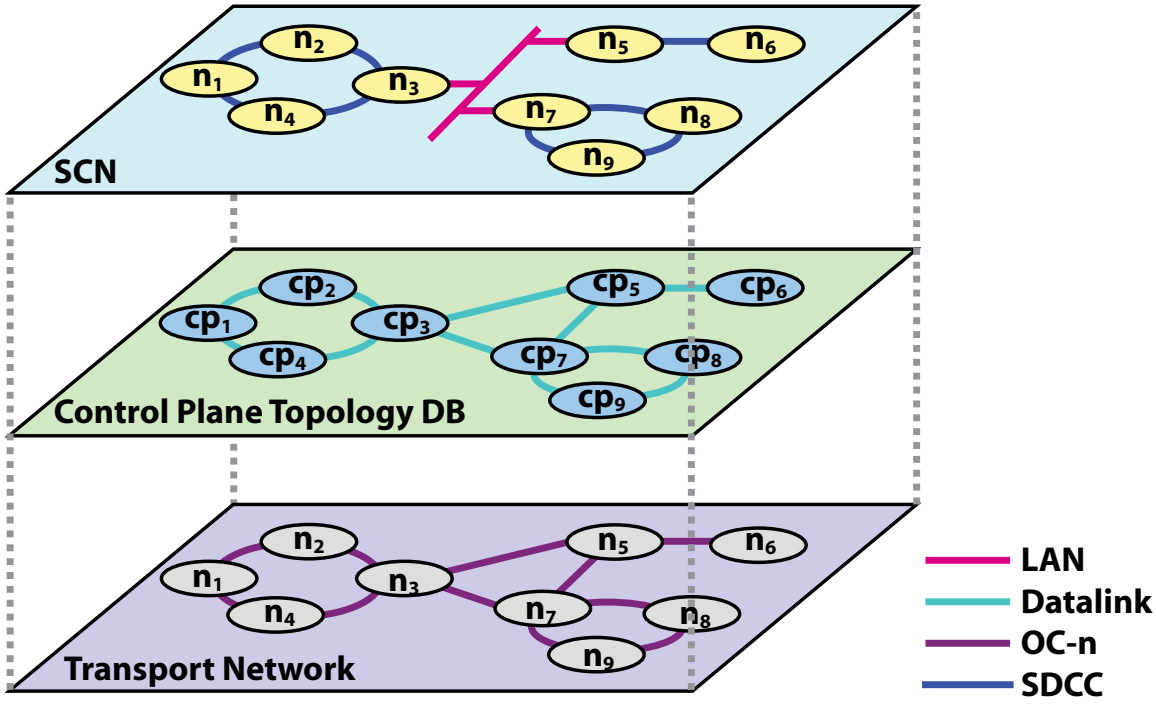

Figure 13: SCN, Control Plane, and Transport Topologies

MCN/SCN Operating Modes

The FLASHWAVE products can support two deployment modes for the SCN. The two deployment modes differ in how the SCN relates to the MCN.

For some customers, ease of provisioning may be their main concern. These customers would prefer not to manage two separate DCNs. A common DCN that supports both the management plane and control plane is sufficient.

With integrated mode, there is only one DCN and the MCN and SCN are merged. The management and control plane applications use the same IP addresses and they share the data communication interfaces (DCC, OSC, LCN, etc.). Figure 14 shows the integrated mode for a SONET node.

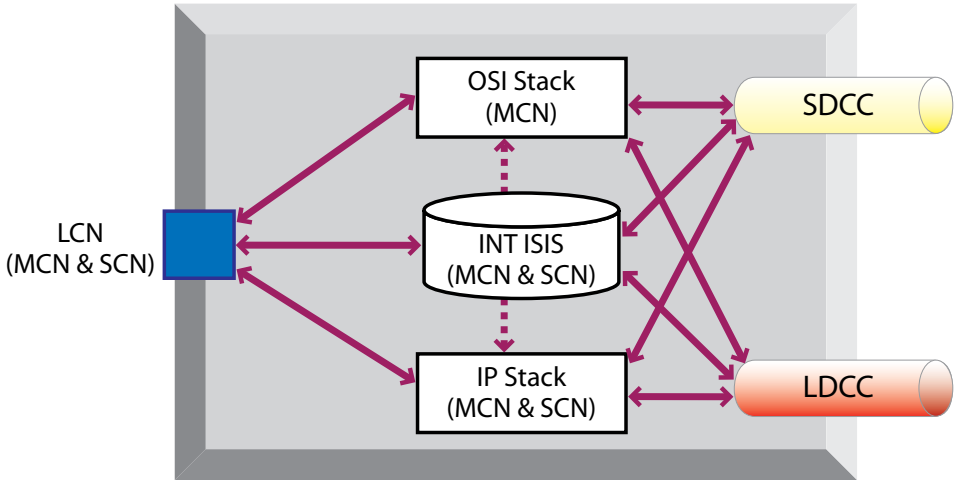

Figure 14: MCN/SCN Integrated Operating Mode

In other cases, customers may want to separate the MCN and SCN for reliability. With separation, failures in one DCN do not affect the operation of the other DCN. Also, traffic volume in the MCN does not affect the SCN and vice versa. Responsibility for address space management can also be divided. The MCN address space can be handled by a traditional IT management group, while network engineering can manage the IP address space assignments for the control plane.

Separated mode has two separate DCNs: one for the MCN and the other for the SCN. The address spaces for the MCN and SCN are fully separated and isolated. It is not possible for a control plane application to reach an MCN IP address, and management applications cannot communicate using the SCN addresses. Figure 15 shows the separated mode for a SONET node.

Further, the data communication interfaces are dedicated to either the MCN or SCN. For instance, the MCN can use the SONET SDCC while the SCN can use the SONET LDCC. For external communications, separate LCN ports are supported on some FLASHWAVE NEs to allow dedicated interfaces to the MCN and SCN. However, some legacy equipment only has a single LCN port. In this case, the MCN and SCN must share the common LCN port and logical separation of the MCN and SCN is necessary.

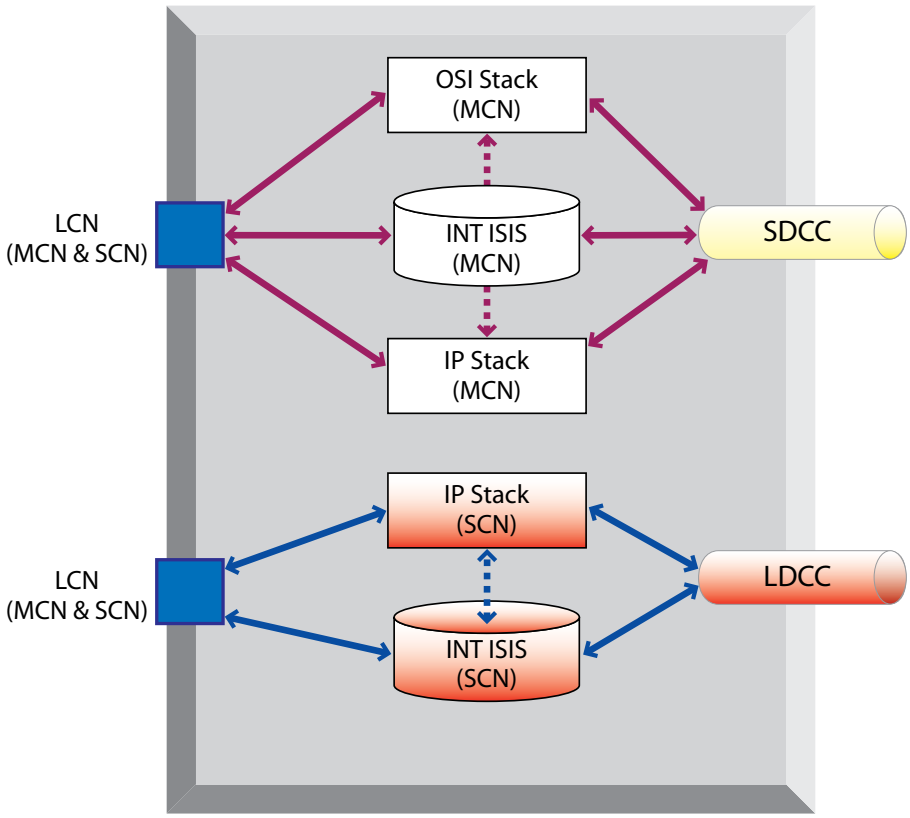

Figure 15: MCN/SCN Separated Operating Mode

Network Applications

There are several principal reasons why carriers consider a control plane for a transport network. Among these are:

- More accurate engineering information
- Faster network provisioning
- Savings through shared network resources for restoration

The significance of these factors should be considered in light of existing capabilities to determine what advantage a control plane can provide.

With the control plane, the network is truly the provisioning database. A circuit can only be reserved and set up if network resources are available, and assignments would not be made to previously assigned resources. With software tools such as NETSMART® 1500 NMS the network is kept in sync with the management system on a continuous basis. Thus if a carrier is already using an appropriate management system, there is much less concern about the accuracy of its database.

Moreover, software like NETSMART 1500 EMS can calculate paths through the network and assign cross-connects in its domain of NEs to establish services. In other cases, a third-party OSS may inventory the NE, circuits and flow assignments through to a provisioning OSS. Once a network is set up to take advantage of these EMS or OSS capabilities, it may be more advantageous to the carrier to continue to use their existing OSS rather than to transition an existing network to the control plane. Today, many existing SONET networks are already fully managed by OSS or EMS, and would not fully benefit by either retrofitting the entire network with a control plane, or by trying to mix a control plane enabled sub-network into the non-control plane-enabled network.

The same may not be true of DWDM networks, which are beginning to be more widely deployed. However, a ROADM network is not fully dynamic as there are often wavelength and connectivity constraints that must be considered with today's DWDM equipment. For this reason, some of the advantages of the control plane cannot be realized until a fully dynamic OADM is cost-effective. Today, lightpaths tend to be put up and left in place, with the OTN network emerging as the switching level in an optical network (ODU switching in an OTU network). Perhaps the transport control plane will play a greater role in emerging OTN networks.

Finally, some carriers are interested in control plane primarily for shared mesh restoration. This is a complex subject in and of itself, and is not addressed in this white paper.

Summary

A control plane-enabled network provides the capability for network elements to automate network provisioning tasks. Control plane functionality enables new dynamic applications such as Bandwidth on Demand. This functionality also reduces network resource requirements through mesh networks and dynamic restoration. It also reduces network operator workload, and promotes interoperability.

This paper provided a high-level description of the Fujitsu approach to control plane architecture for the FLASHWAVE platforms. It covered some general control plane concepts, discussed the control plane applications including discovery, routing, path computation, and signaling, and examined the impacts to the SCN.

Other control plane features such as protection and restoration, multilayer connection provisioning (e.g., Ethernet over VCAT over SONET), and PCE were not discussed in this paper. These topics may be covered in future white papers.

Acronym	Description
ASON	Automatic Switched Optical Network
CCC	Calling Party Call Controller
CR-LDP	Constraint-Based Label Distribution Protocol
CSPF	Constrained Shortest Path First
DCC	Data Communications Channel
DCN	Data Communications Network
DWDM	Dense Wavelength-Division Multiplexing
EMS	Element Management System
E-NNI	External NNI
GCC	General Communications Channel
GMPLS	Generalized MPLS
IETF	Internet Engineering Task Force
I-NNI	Internal NNI
IP	Internet Protocol
ITU-T	International Telecommunication Union-Telecommunications
IS-IS	Intermediate System to Intermediate System
LCN	Local Communications Network
LDCC	Line Data Communications Channel
LMP	Link Management Protocol
MCN	Management Communications Network
MPLS	Multi-Protocol Label Switching
NCC	Network Call Controller
NE	Network Element
NMS	Network Management System
NNI	Network-Network Interface
OADM	Optical Add/Drop Multiplexer

Acronym	Description
OAM&P	Operations, Administration, Maintenance and Provisioning
ODU	Optical Channel Data Unit
OIF	Optical Internetworking Forum
OSC	Optical Supervisory Channel
OSS	Operations Support System
OSPF	Open Shortest Path First
OSPF-TE	OSPF Traffic Engineering
OTN	Optical Transport Network
OTU-n	Optical Channel Transport Unit
PC	Permanent Connection
PCE	Path Computation Element
ROADM	Reconfigurable Optical Add/Drop Multiplexer
RSVP	Resource Reservation Protocol
RSVP-TE	RSVP Traffic Engineering
SC	Switched Connection
SCN	Signaling Communications Network
SDCC	Section Data Communications Channel
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical NETwork
SPC	Soft Permanent Connection
SRLG	Shared Risk Link Group
TE	Traffic Engineering
UNI	User-Network Interface
VCAT	Virtual Concatenation
WDM	Wavelength-Division Multiplexing

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom

References

Architecture

[1] RFC 3945, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," October 2004.

[2] G.8080, "Architecture for the Automatically Switched Optical Network (ASON)," June 2006.

Discovery

[3] RFC 4204, "Link Management Protocol (LMP)," October 2005.

Routing

[4] RFC 2328, "OSPF Version 2," April 1998.

[5] RFC 2370, "The OSPF Opaque LSA Option," July 1998.

[6] RFC 3630, "Traffic Engineering (TE) Extensions to OSPF Version 2," September 2003.

[7] RFC 4203, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)," October 2005.

[8] OIF-ENNI-OSPF-01.0, "External Network-Network Interface (E-NNI) OSPF-Based Routing – 1.0 (Intra-Carrier) Implementation Agreement," January 2007.

Signaling

[9] RFC 2205, "Resource Reservation Protocol (RSVP) – Version 1 Functional Specification," September 1997.

[10] RFC 3209, "RSVP-TE: Extensions to RSVP for LSP Tunnels," December 2001.

[11] RFC 3471, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description," January 2003.

[12] RFC 3473, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol – Traffic Engineering (RSVP-TE) Extensions," January 2003.

[13] OIF-E-NNI-Sig-01.0, "Intra-Carrier E-NNI Signaling Specification," February 2004.

[14] OIF-UNI-01.0-R2-Common, "User Network Interface (UNI) 1.0 Signaling Specification, Release 2: Common Part," February 2004.

[15] OIF-UNI-01.0-R2-RSVP, "RSVP Extensions for User Network Interface (UNI) 1.0 Signaling, Release 2," February 2004.

[16] OIF-UNI-02.0-Common, "User Network Interface (UNI) 2.0 Signaling Specification: Common Part," February 2008.

[17] OIF-UNI-02.0-RSVP, "RSVP Extensions for User Network Interface (UNI) 2.0 Signaling," February 2008.

[18] OIF2005.381.25, "Draft OIF E-NNI Signaling Specification (OIF E-NNI 2.0)," November 2008 [Work In Progress].

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom