

Document Number	EDCS-759846
Revision	2.0
Date	March 05, 2016

ONS Platform Network Connectivity Recommendations

Introduction:

An ONS MSPP, MS TP, and CPT based network is managed over a DCN with LAN and optical DCC/GCC/OSC interfaces. MSTP, MSPP, and CPT DCN interfaces are virtually the same with only some minor differences.

The DCN is primarily used for Management and Control traffic. It is IP based and the packets are carried over Ethernet from the CTC or PRIME Optical to the appropriate GNE (through the LAN interface on the NE's) and from the GNE over DCC (S-DCC/L-DCC) or GCC or OSC to the ENes. For brevity, DCC is used for all these terms throughout the document. Specific names are used when explaining that particular type.

This document provides recommendations and guidelines for ONS Platform networks. The implementation of a DCN is a complex endeavor. A small network, for example, 20 DCC connected nodes, is relatively easy to setup. A large network requires careful thought, both initially and as the network grows. There are a great number of variables and dependencies which influence the design and performance of the DCN.

A given DCN cannot grow indefinitely without suffering management performance issues and at a certain point segmentation of the network into multiple domains may likely be required. If your network is expected to grow beyond 100 DCC connected nodes, plan network design with segmentation in mind. Ad hoc segmentation can prove difficult.

There is no infallible formula that guarantees a specific performance. Each customer's network implementation is different. A number of variables have direct influence over the DCN performance. These include, but may not be limited to the number DCC links in the network, the number DCC terminations on a given node, number of circuits, node main controller type, the network topology, number of GNEs, number of ENes, number of hops to farthest ENE from a GNE, size of a node routing table, OSPF link state database, external routers, external firewalls, simultaneous management users.

The main controller of the NE performs many functions. The DCN connection and DCC routing is just one function. In other words, the shelf controller is not a general purpose router with multiple wire-speed gigabit Ethernet ports dedicated to management traffic. The DCN and DCC routing functions are meant for management traffic of a well-considered DCN/DCC design.

The intended audience for this document includes those familiar with the concepts of Data Communication Network architecture and basic IP configurations. It is intended to supplement the MSTP, MSPP, and CPT documentation.

ONS Network Performance:

The following objectives should be sought in order to optimize network performance:

- Maximize the processing power of the node controller by reducing the traffic load on the DCCs.
- Minimize the CTC-NE traffic flow over the DCCs.
- Understand the types of traffic flow and limit it to ONS management traffic.
- Reduce the number of DCC hops that the traffic flow takes.
- Maximize the DCC link rate, preferring higher rates.
- Optimize the routing that the traffic flow takes.

General:

NEs can be separated into two categories:

- Gateway NEs (GNEs)
- End (or External) NEs (ENEs)

GNEs provide the LAN connection to the NMS for the ENEs. ENEs must go through a GNE to reach an NMS.

In this document, a GNE means a network element (NE) that has LAN interface activated. A GNE can forward or relay traffic between the LAN interface and DCC interfaces.

The term GNE can be used in different contexts when talking about LAN connected NE, TL1 gateways, or with SOCKs Proxy. Unless otherwise noted, the convention used here refers to LAN Connected NEs.

An ENE does not have LAN interface activated. An ENE only forwards traffic between DCC interfaces. An ENE relies on GNEs to communicate with the NMS in the NOC. GNE/ENE ratio is directly proportional to the DCN performance.

With respect to management traffic, an ONS node operates as a Layer 3 router and may provide routing between the LAN and DCC interfaces. ONS node LAN interfaces automatically run Proxy ARP. ONS node LAN interfaces do not run OSPF by default. ONS internal DCC links run OSPF by default.

OSPF on DCCs:

By default, OSPF is enabled on DCCs that is used to communicate reachability of NE in the ONS network. OSPF is also the basis of topology information used by CTC/CPO for complete circuit provisioning.

IP Addressing:

IP Addressing schemes influence optimal packet forwarding.

GNE addressing Rules:

1. All GNEs connected to the same LAN segment should be in the same subnet.
2. All GNEs whose LANs are segmented should have their own subnets.

ENE Addressing Rules:

1. All static ENEs (an ENE were the LAN will not become active) can be on the same subnet as the NMS bound GNEs that are reachable over DCC. This employs Proxy ARP to reach ENEs.
2. For optimal forwarding, ENEs should be on subnets separate from the NMS-bound GNEs. Here, no Proxy ARP is employed to reach ENEs.
3. For proper address summarization, ENE subnets should be assigned in a contiguous fashion to use adjacent subnets. For summarization:
 - Reduces the total number of routing entries in the DCN

- Reduces the impact of link flapping on routing stability
 - Reduces the number of OSPF updates
4. ENE subnets should not cross OSPF Area boundaries.

Network Design:

There are several ways to design the DCN. A simple network may only rely on Proxy ARP. A more complex network may make use of additional Static Routes as well as Proxy ARP. While, in a very complex network, Proxy ARP and Static Routes may be insufficient and you may use OSPF on the LAN.

Design of the DCN must consider the following:

1. The size of the DCC network.
2. The number of GNEs.
3. The GNEs locations in the network.
4. Address planning.

Default Router:

Each NE has a Default Router option which is configurable in CTC. The default router is the router address used when the destination is not in the local subnet and there is no matching entry in the NE routing table. The default router provisioned on an NE results in a default route being installed on the local routing table – destination 0.0.0.0, mask 0.0.0.0, and the gateway is the provided router address. This network route (dest: 0.0.0.0/0 gateway of default router) is called the Default Route.

If the LAN interface is up, the default router, if not 0.0.0.0, is installed on the NE LAN interface and the provisioned router as the gateway. If the LAN is down, then no related route is installed. Default Router is installed only on the local node and not distributed over the OSCs.

If the default router is provisioned as 0.0.0.0, then no default route is installed on the local node (and not distributed over OSC per above).

There is a difference between a static default route and the default router. A static default Route is installed on the local node and redistributed through OSPF on DCC/OSC.

If a non-zero default router is provisioned on a node then when creating a static default Route, the non-zero Default Router must be the static route's gateway.

It is recommended to use a default router of 0.0.0.0 on an ENE or a Stub GNE (Stub GNE is a GNE that does not lead to the NOC over its LAN interface).

When u set the default router to 0.0.0.0 on an ENE plugs into the ENE, no locally generated default route is installed, the ENE will still be visible to the NOC via the GNE.

Proxy ARP:

Proxy ARP is used on ONS GNE to simplify ENE addressing and routing. It is enabled by default on the ONS GNEs unless overridden by provisioning such as enabling SOCKS Proxy. As the ONS network size grows Proxy ARP is not sufficient to meet performance needs. Forethought should be put into designing the network for growth. There are numerous considerations in migrating to large, balanced optimal DCN and ONS networks. One goal would be to employ a dynamic routing protocol on the DCN with OSPF on the LAN and the optimally deployed DCC network(s).Proxy ARP guideline are as follows:

1. All nodes are in the same IP subnet.
2. Each ENE should be within 5 SDCC (or OC3-OSC) hops and 8 LDCC (or FE/GE OSC & high rate GCC) hops from the GNE.
3. There can be only one Proxy ARP Server (PAS) in an IP subnet.
4. When there are multiple GNEs on an IP subnet the GNE with the largest MAC address becomes the PAS.

5. DCC connected peers must be in the same OSPF Area as the PAS GNE, or share one common Area if the PAS GNE is in more than one Area.
6. No Load balancing is supported between the GNEs.
7. Proxy ARP only works well with small networks (less than 20 NE's); in larger networks this can result in sub-optimal routing from the DCN to DCC and degrades performance.
8. Not recommended architecture for large and complex network build outs.

Static Routes:

With Proxy ARP, there is only one PAS which can lead to sub-optimal routing. Static routes can be employed to remedy this inadequacy.

Static routes may also be used on the customer's LAN to manage paths to various ONS nodes.

As the size of the network increases beyond the limits of PAS, it would be helpful to employ static routes. Ultimately, dynamic routing such as OSPF is a preferred solution.

As the name implies, static routes are static and as networks grow or change, static routes may need to be updated.

Static routes are not ideal in a network that needs to track link failures and to adapt traffic flows to failures.

1. Static routing is used when OSPF is not running on the NE LAN interface.
2. Static routes may also be used on the DCN side to manage paths to various ONS nodes.
3. Static routes on the GNE tell the GNE how to reach a destination and to advertise the routes to the DCC network.
4. When there are multiple GNEs, the same static route needs to be provisioned on all GNEs so that the ENE can determine the best route.
5. If the destination is on the same subnet, no static routes are needed because the subnet route is advertised automatically by the GNE (note, this can get quite complex when stub GNEs are advertising the same subnet route).
6. If the destination is on the same subnet but the LAN is segmented, static routes can be used to maintain visibility.
7. When there are multiple GNEs, special attention must be given to which GNE is to be the gateway to a certain ENE. The simple design rule is to evenly distribute ENEs to each GNE such that each GNE covers an equal number of ENEs.
8. As the size of the network increases, it may become helpful to employ static routes. Dynamic routing such as OSPF is a preferred solution.
9. A static default route (with 0.0.0.0 as the IP and mask as 0.0.0.0) provisioned on the GNE will add a route to all DCN subnets on the non-LAN connected nodes.

Dynamic Routing:

To allow for dynamic exchange of reachability information between NEs and the NOC, OSPF can be enabled on the LAN interface of a GNE and on the links of routers connecting to the GNEs.

OSPF on LAN also provides for load balancing among multiple GNEs, and routing redundancy between the DCN and DCC networks.

OSPF on DCC and LAN:

Guidelines:

1. ENE congestion: Each ENE should be within 5 SDCC (or OC3-OSC) hops and 8 LDCC (or FE/GE OSC) hops from the GNE. Maximum 8-hops limitation means that you can have 16 ENEs between two GNEs so that when both GNEs are available, the ENE has nearest GNE within 8 hops. When one of the GNE fails, there will be a congestion on DCN but the network would still continue to work with slower response.
2. GNE congestion: In general, each GNE should not be routing for more than 20 ENEs at which point additional GNEs should be considered.
3. DCC OSPF Area congestion: Each DCC OSPF Area should have less than 50 nodes. This estimate is normalized across ONS 15600, 15454, 15310 and ONS 15454 MSTP.
4. The position and number of GNEs in the network is crucial, violation of these first three rules can produce slow performance of NMS and long delays.
5. All the nodes in a give OSPF Area (including Area 0.0.0.0) must be contiguous. The results of segmented OSPF Areas are undetermined and may include instability.
6. Initial planning of the network should stress addressing of the DCC network so that OSPF summarization is possible and optimized as the network grows. It is highly desirable to make a node with multiple DCC OSPF Areas a GNE.
7. It is recommended to assign different subnet addresses to different Areas.
8. The router priority defaults to 0 so that a GNE will not be elected as the OSPF Designated Router (DR); unless there is an absolute need, this should be left as the default.
9. Avoid the use of external device access using DCC. Use the DCC bandwidth for ONS management only. If external devices are connected to ONS boxes, this adds to the route table size, the number of LSAs in the network, and the size of the LSDB. DCC bandwidth is limited.
10. Avoid plugging in a hub or switch or third party devices to the LAN port of an ENE. This will add to the route table and the OSPF LSDB. Limit only to temporary use of a technician's PC to manage the NE using the craft port.
11. SOCKS proxy settings will terminate all the CTC/CTM connections on the GNE node. This increases the packet processing on the GNE. So there should be sufficient GNE nodes if the network is based on SOCKS proxy settings.
12. OSPF provides dynamic routing, load balancing among multiple GNEs, and routing redundancy access to the DCC network.
13. OPSF is always enabled on the DCC networks and can be enabled on the LAN.
14. By default if other DCC links exist in OSPF Area 0.0.0.0, new DCC links are assigned to OSPF Area 0.0.0.0 – the backbone OSPF Area.
15. Avoid DCC OSPF Area segmentation or churn:
 - a. Prior to DCC link turn up, its OSPF Area ID needs to be identified.
 - b. A new DCC turn up needs to be assigned its identified OSPF Area at turn up.
16. If all existing DCC links are in numbered Areas then a new numbered Area DCC link will be assigned by the NE. This new DCC links should be verified to be in the desired OSPF Area.
17. If OSPF is enabled on the LAN, then the DCC links must be in a numbered OSPF Area.
18. All BSLR links must be in the same OSPF Area.

19. Provisioning of OSPF parameters must match on both sides of the link. When changing DCC metrics, ensure that SDCC metric value is 3 times the metric value for LDCC.
20. When OSPF on the LAN is enabled, static routes through the LAN interface, including the default gateway settings are automatically disabled.
21. The following rules apply to Area IDs on the LAN interface:
 - a. Both the router interface and the LAN interface must be set to the same Area ID.
 - b. The DCC interfaces must be on a different Area than the LAN interface
 - c. The DCC interfaces cannot be on Area 0.0.0.0.
22. OSPF Timer, Router Priority, Authentication, Metric parameters must be set the same on the LAN interface and the router interface.
23. Range tables summarize the OSPF advertisements of reachable IP addresses of an Area to reduce the number of LSAs.
24. Virtual Links are needed when LAN Area is not 0.0.0.0; it provides a virtual link to the backbone Area. OSPF Virtual Links should be used only under extremely special cases. OSPF Virtual Links incur a performance cost.
25. Area Border Routers (ABRs) are recommended to have a connection to the backbone Area.
26. Keep the backbone Area contiguous and simple.
27. Maintain an optimal Area size to contain LSA flooding scope, generally to a desirable size of less than 50 NE's.
28. Plan Area ID assignment to allow for network expansion.
29. Keep a subnet within a single Area.
30. Ensure there are redundant Area 0.0.0.0 links and Inter-Area links.
31. Ensure that OSPF is properly configured on the routers attached to the LAN before OSPF is enabled on GNEs.
32. Balance summarization with forwarding optimization; addressing affects summarization and traffic forwarding.
33. It is recommended to limit total entries in the routing table to fewer than 500 routes. Various system and network resource over use can lead to unpredictable routing problems; routing table size is one that is easiest to measure.
34. Avoid or at least be conservative when it comes to modifying OSPF default settings on GNEs.
35. Maximum number of OSPF Areas that can be supported on an NE is 5. There is no hard limit, but as Areas/LSDB grows, so does the CPU usage.

OSPF on the LAN

In order to provide dynamic routing with the DCN, OSPF can be enabled on the GNE LAN interface. In addition to the traditional OSPF requirements, the ONS has some additional rules for how to use OSPF. With OSPF on the LAN, the following points should be noted:

- **Static Routes:** All static routes and default routes on the GNE are disabled when OSPF is enabled on the LAN.
- **Area ID:** The area ID on the LAN cannot be the same as any of the area IDs on the DCC links. CTC will not allow any new DCC links to be in area 0.
- **Link metrics:** The default link metric for the LAN interface is 10. Link metrics can be changed, but this should be done carefully and consistently.

- **Virtual Links:** One Virtual Link is allowed per node. It is supported only over the LAN interface (because Virtual Links require a numbered interface). The neighbor ID is the neighbor's Router ID and not the interface address.
- **OSPF Options:** Only the E1 option is supported. This means that the Area supports external routing. An option E mismatch between neighbors will prevent the OSPF LAN session from coming up. Note that the E1 option is not configurable on the ONS.
- **Redistributing external routes into the ONS:** DCN based external routes are communicated with LSA type 5. The default route is described by an LSA type 5. An upstream router can redistribute default routes in the same way as the IOS originates default information.

OSPF Area Border Routers:

While designing an ONS network, one needs to pay attention to the Area Border Routers (ABRs). An ABR connects one or more regular OSPF Area to the backbone Area. The following should be considered:

1. Without OSPF Summarization turned on:
 - a. Each NE's 32-bit host route is being advertised by every ABR as LSA Type-3s.
 - b. If an ABR has more than one numbered Areas, each numbered Area will advertise each NE via LSA Type-3.
 - c. Area zero will get the concentration of all these LSA Type-3s.
 - d. If OSPF summarization is not maintained it is possible to produce LSA storms. LSA storms are typically denoted by rolling HELLO FAIL alarms and possibly high NE CPU usage which can lead to temporary loss of visibility.
2. Network and NE performance may be degraded, one should consider incrementally implementing OSPF summarization as OSPF Areas are added to the network

OSPF Summarization:



Note

Intra-area routes refer to updates that are passed within the area. Inter-area routes refer to updates that are passed between areas. External routes refer to updates passed from another routing protocol into the OSPF domain by the Autonomous System Border Router (ASBR).

An inter-area summary route is also called an area range. A summary route is created when OSPF is first enabled on the LAN (range table entry with the GNE's local subnet, area, prefix length, and advertisement disabled). As a result, the NE addresses in the local subnet are suppressed by the GNE ABR (this automatically created summary is not advertised into area 0), with the exception of the ABR's own address. If the NE address or the subnet is changed later, the previous summary entry is not automatically updated, which might lead to undesirable results in routing. Thus, after OSPF LAN interface changes, Area changes, and address changes, OSPF range table entries should be checked for correctness.

Under a non-zero area, three summary LSAs are advertised:

- A Type 3 for the local subnet
- A Type 3 for itself with a /32 length
- A Type 4 for the attached Router

In the absence of range table entries, the GNE advertises a Type 3 LSA for each ENE address with the prefix length of 32 (a host route) to area 0. As a result, all ENEs in the numbered area are reachable via inter-area host routes on the Router.

**Note**

An incorrect range table entry can result in the GNE advertising a Type 3 LSA for each ENE behind it with a prefix of 32 in to the backbone Area, plus adding invalid Type 3 LSAs present in the range table.

Inter-Area Summarization:

ONS ABRs have the ability to do address summarization towards area 0.0.0.0 or towards ALL the attached areas.

When creating an address summary range table entry you specify the address range, the source area, and the range prefix (mask) and whether to advertise or to suppress advertisement. The most common form of a range table entry is summarizing an attached area towards the backbone.

When OSPF is enabled on the LAN, a default range table entry is created. This default range table entry suppresses all the addresses in the GNE's subnet Except it's own address. Its own address is advertised as an LSA Type 1 and is installed as a host route (/32) on the attached router. The default range table entry specifies the GNE's subnet, the attached area, the GNE's network mask length and has the advertisement disabled (un-checked). Note that there is a restriction on ONS range table entries. You cannot have two summaries with the same network number and with the same prefix length or larger. For example, if you create a summary for 10.11.12.0/24, you cannot create one for 10.11.12.0/25 through /32.

When the automatic range table entry is present; the GNE suppresses all addresses in the GNE subnet towards the attached router. The router has a host route to the GNE in its routing table. If the router, the GNE, and the ENEs are all in the same subnet, the Proxy ARP Server GNE handles any traffic bound for the ENEs.

If the automatic range table entry is removed, then all the host addresses that are in the same subnet will be advertised in the OSPF. The advantage here is that the router, through OSPF, learns all the NEs which include a cost so that the router can route to the NEs more effectively. Basically, in this case it would defeat the PAS. The disadvantage is that there are more routes in the router and associated ONS NEs.

Routing becomes more complicated when the NE uses different subnets. Complicated and improper summarization leads to suboptimal forwarding.

The second form of summarization involves suppressing advertisement. A common range table entry that suppresses advertisement other than the automatic range table entry, involves suppressing summaries from the backbone are into the attached areas on an ABR.

External Devices:

External Devices off of ENEs are highly discouraged and should be avoided. When an external device is connected to an ENE it impacts the following resources which degrade performance:

- Adds a host route.
- Adds a network route.
- May require a static route(s), which adds more routes.
- Introduces foreign traffic on the DCC network.
- On an ABR, each external device will cause an LSA-4 to be created and advertised across OSPF Area boundaries.

Inconsistencies with OSPF on the LAN Provisioning:

When a DCC network has OSPF enabled on the LAN, all the GNEs should have OSPF enabled on the LAN. If one or more of the GNEs do not have OSPF enabled on the LAN while the rest of the network does, errors may occur. If an OSPF enabled LAN network has a GNE that does not have OSPF enabled on the LAN, the GNE becomes an Autonomous System Border Router. In OSPF, the GNE causes Type 4 and 5 LSAs to be generated throughout the regular Areas. This distorts the routing table on the GNE and cause visibility issues.

ONS Communication Channels:

- DCC
 - Data Communication Channel (DCC)
 - DCC carries management channel communication for a TDM network
 - DCCs are routable and can be used to discover
 - ONS network topologies using Open Shortest Path First (OSPF)
 - DCCs interfaces appear in the IP routing tables as “pdcc”
- OSC
 - Optical Service Channel (OSC)
 - OSC is a bidirectional channel connecting two adjacent nodes in a DWDM ring
 - An OSC signal uses the 1510-nm wavelength and does not affect client traffic
 - The primary purpose of this channel is to carry clock synchronization and management channel communications for the DWDM network
 - OSCs are routable and can be used to discover
 - ONS network topologies using Open Shortest Path First (OSPF)
 - OSCs interfaces appear in the IP routing tables as “pdcc”
- GCC
 - Generic Communication Channel (GCC)
 - Some line cards use a digital wrapper function (ITU-T G.709 compliant) and formats the DWDM wavelength so that it can be used to set up GCCs for data communications, enable FEC, or facilitate performance monitoring
 - GCCs are routable and can be used to discover
 - ONS network topologies using Open Shortest Path First (OSPF)
 - GCCs interfaces appear in the IP routing tables as “pdcc”
- PPC
 - Provisionable Patch Cord (PPC)
 - PPCs are required when the TXP, MXP, GE_XP, 10GE_XP, ADM-10G, or ITU-T line card is installed in a different node than the OCH filter ports
 - They can also be used to create OTS-to-OTS links between nodes that do not have OSC connectivity
 - PPCs between TXP, MXP, GE_XP, 10GE_XP, ADM-10G, or ITU-T line cards and OCH filter ports are not routable
 - Only OTS-to-OTS PPCs are discovered by OSPF and make adjacencies and can be used to discover ONS network topologies

- PPCs between TXP, MXP, GE_XP, 10GE_XP, ADM-10G, or ITU-T line cards and OCH filter ports generate only OSPF opaque LSAs to be used to discover ONS network topologies
- PPC interfaces are not included in the IP routing tables
- GCCs and DCCs are not required for PPC creation

OSC Types:

The OSC is a bidirectional channel connecting all the MSTP NEs (such as Terminal, OLA, & ROAM) that is able to carries general purpose information without affecting the user transport traffic.

The OSC is a single optical wavelength. This wavelength runs on the same fiber carrying the optical payload wavelengths but is located out of the DWDM band (1510 nm) of the optical amplifiers.

M12, M6/M2 Chassis, and OSC rates:

The OSC optical rate supported on a given MSTP chassis is predicated by the type of OSC card supported by the chassis.

M12:

M12 OSC is supported by the following two fixed optic cards:

- OSCM
- OSC-CSM

Both the OSCM and OSC-CSM provide a supervisory data channel at the OC3 SDCC fixed rate of 192 kbps. Both cards also provide a 100 Mbps User Data Channel (UDC) that run on the OC3 payload.

M6/M2:

The M6 and M2 OSC is supported by the following two cards:

- TNC
- OSC-CSM

OSC, UDC, and VoIP are supported by two SFPs on the TNC card. UDC and VoIP are supported only on the M6, the M2 ECU does not provide UDC or VoIP connections.

- SFP0 Rates: OC3, FE, and GE
- SFP1 Rates: FE and GE

The OSC-CSM card is also supported on the M6 and M2 providing fixed OC3 rate DCC (192 kbps) and UDC (100 Mbps).

TNC provides a supervisory data channel for node-to-node communications in three forms: OC3 SDCC at 192 kbps, FE and GE. TNC provides 100 Mbps UDC. TNC provides 10 Mbps VoIP channel.

ETHERNET OSC:

The TNC supports high-speed OSC running over an underlying Ethernet link which is called an Ethernet OSC. The high-speed OSC can be used when the TNC SFP is either 100Base-FX or 1000Base-X.

On TNC, the same physical link can carry three types of traffic:

- OSC

- VOIP
- UDC

The OSC, VoIP and the UDC are separated by three different VLANs. These VLANs are not configurable. The Ethernet-OSC is a supervisory channel between the two Cisco M6 and M2 nodes. Each TNC card supports up to 2 Ethernet-OSCs. Each M6 shelf, which can have two TNCs, can support up to 4 Ethernet-OSCs. Each M2 shelf can have one TNC, can support up to two Ethernet OSCs.

It is possible to create Ethernet-OSCs on the node controller shelf as well as subtended M6 shelves. In one multi shelf node, the total number of Ethernet-OSCs depend on the “degree” of the node.

Provisioning of Ethernet-OSCs is supported via both CTC and TL1. All OSC attributes supported on SDCC-based OSCs are supported on Ethernet-OSCs too. This includes OSPF link-metric for Ethernet-OSCs (just like for SDCCs and LDCCs).

OC3 SDCC-based OSC is implemented using the SDCC bytes of an OC3 signal. Ethernet-OSCs run on an Ethernet link. The internal hardware path to the TNC CPU for the Ethernet-OSC is completely different from the OC3 OSC. It uses the BCM Ethernet switches present on the TNC board. The OC3 OSC hardware path to the TNC CPU use FPGAs.

FE vs. GE OSC:

The main difference between the FE and GE OSC is basically the bandwidth and the SFP reach (optical module):

- For FE, the SFP reach is 44dB.
- For GE, the SFP reach is 20dB.

No timing is required on FE and GE. No issues related to TNC active and stand-by timing is present as the OSC FE/GE traffic is terminated on L2 switches in TNCs and the internal communication between the TNCs is Ethernet and not SCL (TDM frames).

FE OSC share 100 Mbps supervisory data channel and 100 Mbps Mbps UDC (and VoIP). The supervisory data channel is preference over UDC and VoIP on FE OSC, so that the supervisory data channel can eliminate UDC. This is not a concern with GE OSC, the UDC gets 100 Mbps and the rest belongs to the supervisory data channel.

The best-case performance of these OSC types is captured in the table below. The testing was done on a simple back-to-back (adjacent node) configuration. The difference in Measured vs. Theoretical bandwidth is because of throttling by the controller CPU.

OSC Type	Theoretical bandwidth	Measured bandwidth
OC3	192Kbps	166.93kbps
FE	100Mbps	721.57kbps
GE	1Gbps	771.34kbps

DCC GNE to ENE Hops:

Consider the guideline of 5 hops to reach a GNE over a 192 kbps SDCC or 8 hops to reach the GNE over a 576 kbps LDCC. For OC3 OSC the number of hops should be no more than 5. For FE/GE OSC the number of hops should be no more than 8. GCC can have configurable rates; for 192 kbps GCC the number of hops should be no more than 5 and for any higher rates should be no more than 8.

These rules are geared toward LAN access and are based on ideal conditions. Consider the difference between SDCC and LDCC. Three times the bandwidth of the LDCC does not mean three times the hops achieved. Consider the aggregation effect at each hop as traffic flows accumulate at the egress points. The amount of DCC traffic is greatly increased with the addition of traffic from excessive aggregation. Consideration should also be made for the higher arrival rate effect of LDCC, Ethernet OSC, or high rate GCC and processor impact.

GCC:

A DWDM based GCC is analogous to a TDM SDCC or LDCC. A GCC is associated with a line card trunk port. The data carried over a GCC is formatted in a digital wrapper defined by ITU-T G.709.

There are legacy GCC line cards that only support a 192 Kbps rate and there are high speed GCC line cards and CPT ports that allow a provisionable rate.

A maximum of five GCC channels can be created on the Cisco ONS 15454 M12 shelf and ten GCC channels can be created on the Cisco ONS 15454 M2 or Cisco ONS 15454 M6. High speed GCC can be created when both the NE and appropriate line cards are in Cisco ONS 15454 M2 or Cisco ONS 15454 M6 shelves. The legacy 192 kbps GCC can be selected when both sides are Cisco ONS 15454 M12 or one side is a Cisco ONS 15454 M12 the other side is a ONS 15454 M2 or Cisco ONS 15454 M6 shelf.

DCC limits:

There are maximum numbers of DCC allocation per NE type. These are provisionable limits beyond which the software does not allow. However, in practice, other factors may lower the useful limit to less than the absolute provisionable limit.

A number of factors influence the number of DCCs on an NE. These include, but may not be limited to the number of DCC links in the network, the number DCC terminations on a given node, number of circuits, node main controller type, the network topology, the number of GNEs, the number of ENEs, hops to farthest ENE from GNE, size of node routing table, OSPF link state database, external routers, external firewalls, simultaneous management users, number of links in an DCC OSPF AS, number of DCC terminations per NE, the number of external devices, if the node is an ABR, how summarization is used, and the pros and cons of using the routing table as an indication of network stability.

SONET BLSR Considerations:

- The OSPF Area ID must be the same for the DCC links in a BLSR.
- BLSR information is propagated in special OSPF packets called an Opaque LSA.
- Opaque LSAs do not cross OSPF Area boundaries.
- When changing OSPF Areas of a BLSR, both “East” and “West” ports of a BLSR on a node must be changed at the same time.
- Since a BLSR has a maximum limit of nodes and a BLSR must reside in a single OSPF Area, the guidelines for GNE placement must be considered.

- 32 Ring Two-fiber BLSR
 - With Release 5.0 the ONS 15600 can support 32 concurrent two-fiber bidirectional line switched rings (BLSRs).
 - Each BLSR can support up to 24 ONS nodes.
 - Because the working and protect bandwidths must be equal, you can create only OC-12, OC-48, or OC-192 BLSRs.

SOCKs Proxy:

SOCKs proxy is an internet protocol that provides routing of packets between a client and server applications by means of a proxy server. SOCKs proxy allows external clients outside a firewall to connect to servers inside the firewall. SOCKs proxy also allows client behind a firewall to connect to a SOCKs server in order to have access to external services through the firewall.

SOCKs proxy adds additional burden on the GNE node. Sockets are terminated and packets are rewritten by the GNE node which adds to the controller card load on the GNE. There should be sufficient number of GNE nodes so that an optimal performance is achieved.

CTC and SOCKs Proxy:

SOCK proxy is an optional feature that provides ease of NE reachability and ONS related security options. CTC and CPO are SOCKs proxy aware applications.

SOCKs proxy is a standard-based protocol to allow a gateway to act as a TCP/UDP relay or proxy between two end hosts. SOCKs proxy uses the typical client/server model where the end user host (NMS user) is a client and the GNE is the server with a given ENE as the final destination.

The CTC will poll the GNE for proxy information. If SOCKs proxy is enabled on the GNE then the CTC learns all the ENEs being served by the GNE. The CTC then attempts to set up one connection to each ENE. The GNE maintains a translation table for each connection between the CTC and the respective ENE. With CTC using SOCKs proxy, direct reachability to an ENE is not required.

SOCKs proxy terminates all the CTC/CPO connections on the GNE node. This increases the packet processing on the GNE. So, there should be sufficient number of GNE nodes if the network is based on SOCKs proxy settings. The CTC/CPO implementation of SOCKs proxy has a built in GNE load balancing mechanism. SOCKs GNE/ENE ratio determines the speed of network discovery.

SOCKs proxy provides the following functionality:

- SOCKs proxy is a standard protocol that enables a gateway device to act as a TCP/UDP relay (proxy) between two end hosts.
- The TCP socket port 1080 is used to poll for SOCK proxy control information from the NEs.
- A GNE can be configured to enable a firewall; this allows the NMS to manage the GNE and all the DCC connected NEs behind it but the NMS does not have direct IP connectivity to the DCC connected NEs.
- Craft Access Only: With the firewall enabled, any CTC connected to the front panel craft port of an ENE is isolated to the DCC network and prevents access to the LAN connected to the GNE.

There are three options available to enable SOCKs proxy:

- GNE – Gateway Network Element: LAN + SOCKs Proxy + Enable Firewall.
- ENE – External Network Element: SOCKs proxy + Enable Firewall + Craft Access Only.

- SOCKS proxy Only: SOCK proxy.

Other features associated with SOCKs Proxy:

- ENEs can use private or non-routable address.
- ENEs are not being served by a PAS.
- A GNE can only do SOCKs proxy for ENEs in its own OSPF Area(s).
 - When multiple DCC OSPF Areas are employed, each Area should have its own SOCKs proxy GNE.
- A SOCKs proxy GNE does not support OSPF on the LAN.
 - Static routes may need to be provisioned on the GNE to reach an NMS.
 - When SOCKs proxy is enabled these static routes are not advertised into the DCC network.
 - DCN routers do not need static routes to the ENEs since all traffic is forwarded by the GNE.

Multi-Shelf Nodes:

Multi-shelf nodes consist of a Node Controller shelf and a number of subtending shelves. The Node Controller shelf is the center of the inter-shelf communications. There is no inter-subtending shelf communications, that is, no sub-shelf communicates directly with another sub shelf.

DCCs terminate on the Node Controller. If a DCC is associated with a subtending shelf, it must use inter-shelf communications to terminate on the Node Controller. Minimizing DCC traffic flow, benefits the inter-shelf communication resources.

M12 node controller based multi-shelf systems require a special note with respect to DCN routing. An M6 node controller based multi-shelf systems have an integrated multi-shelf switch. This is applicable for M6 node control only if an external multi-shelf switch is used with an M6 node controller.

M12 node controller based multi-shelf systems use a dedicated switches for inter-shelf communications. These multi-shelf switches can take the form of external switches or MS-ISC cards. Both, management (public) and inter-shelf (private) communications are handled by the multi-shelf switch. The node controller LAN interface connects to the switch forwarding both public and private traffic. The multi-shelf switch uses VLANs to separate the public and private traffic. Public traffic is forwarded to your DCN and the private traffic is forwarded to the sub-shelves.

As noted above, various routes are installed on a node based on the state of the LAN interface. In the M12 based multi-shelf, your DCN interface can be down but the TCC LAN interface can remain up. Based on the NE's role in the network, this poses two areas of concern.

M12 Multi-Shelf GNE: In the scenario where the DCN facing interface on the multi-shelf switch goes down but the TCC LAN interface remains up, the nodes routing table can continue to forward packets out the TCC LAN interface. Since the DCN connection on the switch is not up, the DCN bound traffic can be forwarded to the multi-shelf switch but get discarded at the switch since the multi-shelf DCN interface is down. This can result in the loss of visibility to the GNE and some of its ENEs.

M12 multi shelf ENE: The default router for an ENE should be provisioned as 0.0.0.0. This prevents a default route from being installed on the ENE. Note that an M12 Multi-shelf ENE has its LAN interface up because it is connected to the Multi-shelf switch. Because this LAN interface is up, the NE's local subnet route is installed. This poses a minor problem when craft PCs or other EMS hosts us the same subnet as the ENE. This is similar to the case above in which the host's EMS traffic is forwarded to the multi-shelf switch and discarded causing visibility issues.

Secure Mode:

An NEs LAN port can be generalized as having two modes; repeater mode and secure mode. In repeater mode all LAN ports share the same LAN interface. In secure mode there are two LAN interfaces, the back interface and the front interface.

- The back interface consists of either the backplane port (ANSI) or the FMEC port (SDH) on an M12 shelf. On the M6, both interfaces are on the ECU; the back interface is the EMS port and the front interface is the CRAFT port.
- The front interface consists of the front panel LAN ports on the shelf controllers or the CRAFT port on the M6 ECU.
- The back and front interfaces are assigned separate IP addresses.
- The back interface is called the private interface and is intended to provide a private LAN connection to the NMS.
- The front interface is called the public interface and is intended for craft access for field technicians.
- This allows a well known common address for craft access and isolates access from the NMS by the field technicians.
- Public and secure interface IP addresses must reside on different IP subnets. PAS is disabled.
- SOCKS proxy is used and is automatically enabled when secure mode is provisioned.
- OSPF or RIP cannot be enabled on either LAN interface.
- Enabling secure mode results in a TCC/TNC reset.
- Secure Mode can be provisioned to be ENABLED and UNLOCKED or ENABLED and LOCKED. When Secure Mode is provisioned to be LOCKED, it is irreversible.

Glossary

ABR: OSPF Area Border Router; a node with multiple Areas and connection to the backbone Area 0)

ARP: Address Resolution Protocol

AS: OSPF Autonomous System

DCN: Data Communication Network

DCC: Overhead bytes on SONET frame used for management traffic (SDCC and LDCC)

ENE: External NE; a node without DCN LAN connection and only DCC connection, sometimes also referred to as node with SOCKS proxy and ENE enabled

GCC: Generic Communication Channel

GNE: Gateway NE; a node with DCN LAN connection and with SOCKS proxy enabled

LSA: Link State Advertisement of OSPF

LSDB: Link State Database

NMS: Network Management System

NOC: Network Operation Center

OSC: Optical Service Channel

OSPF: Open Shortest Path First

PPC: Provisional Patch Cord

RIP: Routing Information Protocol

References

Zhang, R. Optical Networking Systems IP Management Solutions, Cisco Press, 2007.

