Cisco 7600 Series
Solution and Design Guide - Metro Aggregation

The Cisco 7600 Series Internet Routers is a solution for service providers, offering high-performance optical networking with uncompromising line-rate packet services. This document addresses one of the key applications of Cisco 7600 Series in the service provider market: metro aggregation.

Metro Aggregation Market

Metro aggregation refers to the aggregation of access connections in a metropolitan area network (MAN). Metro networks, as the name implies, are found in metropolitan areas characterized by a high concentration of business and commercial facilities in a relatively small geographical area. A high concentration of customers prompts economies of scale, enabling service providers to offer high-bandwidth network services at relatively low costs. The recent advent of optical transport technologies has fueled the growth of metro network services, as the increased span of optical fiber (up to tens of miles) eliminates the need for optical repeaters, and makes it possible to provision dark fiber in quantity to support network services within metro areas.

Traditional metro network services have been based primarily on SONET and ATM technologies to offer high-bandwidth connectivity. Service providers build SONET and ATM networks to interconnect customers' locations within metro areas, and to provide Internet access and long-haul connectivity to other metro areas. Customers will access such services via SONET and ATM access links from their premises to service providers' points of presence (POPs), and the customer edge equipment has to support WAN interfaces such as SONET and ATM. Customer edge equipment can be owned by the customer or by the service provider.

An emerging alternative to the traditional SONET/ATM-based metro networks is the use of Gigabit Ethernet (GE) as the access link between customers' premises and service providers' POPs. Service providers are beginning to offer GE-based services over fibers within metro areas from customers' premises to service providers' POPs. The advantages are a simpler technology, higher bandwidth, and less-expensive edge equipment—a GE interface instead of a SONET or ATM interface. Customers can use these GE services to access the Internet or to connect to their other locations either within the same metro area or over long distance using the service provider's long-haul backbone. With the right equipment such as the Cisco 7600, a network can support various service features. For example, service providers can offer "fractional GE" services through traffic shaping on the GE interfaces, when there is not enough customer traffic to justify a full GE. Another service feature is to use virtual LAN (VLAN) technology to support multiple customers from the same location over a single GE link. In this case, there is a GE switch on customer's premises supporting multiple VLANs, and the GE access link is used as a VLAN trunk.

This document addresses the various aspects of designing a metro aggregation solution using the Cisco 7600. The focus is on service providers' POPs. It first provides a configuration overview of the Cisco 7600 to describe the product components and configuration guidelines. Then it presents the Cisco 7600 solution in metro aggregation, starting from the aggregation interfaces, through details of interconnecting multiple aggregation routers in the POP and connecting the aggregation layer to server farms and core backbone routers, to traffic engineering and guidelines for sizing uplinks. Technical details and reasoning for the architecture are provided, so that readers can adapt the architecture components and design techniques to meet their customers' networking requirements.

Cisco 7600 Configuration Overview

The Cisco 7600 is an integrated router solution leveraging the Catalyst® 6500 chassis and enhanced core components, together with a new set of high-performance WAN linecards called Optical Services Modules (OSMs). The following is a summary of the Cisco 7600 system components and configuration options.

Chassis

There are three chassis available: a six-slot chassis and a nine-slot chassis both with horizontal slots, and a nine-slot chassis with vertical slots. The six-slot chassis provides six horizontal slots with side-to-side airflow and redundant power supplies. The slots are numbered 1 to 6 from top to bottom. The nine-slot chassis provides nine horizontal slots with side-to-side airflow and redundant power supplies. The slots are numbered 1 to 9 from top to bottom. The other nine-slot chassis provides nine vertical slots with front to back airflow and redundant power supplies. The slots are numbered 1 to 9 from right to left. In all three chassis, slot 1 is for the Supervisor Engine 2 (SUP2), while all other slots can support all linecards. Optionally, slot 2 can support a redundant SUP2, slot 5 can support a Crossbar Switch Fabric Module, and slot 6 can support a redundant Crossbar Switch Fabric Module.

Core Components

- The Supervisor Engine 2 (SUP2) is the processor card, and it has two built-in GE switch ports. An optional redundant SUP2 is supported, and the two GE ports on the redundant SUP2 can be active all the time.
- The Policy Feature Card 2 (PFC2) is a daughter card on SUP2. It is the layer 3 hardware-switching engine, and it also provides hardware-based access-list checking, QoS classification, and input policing.
- The Multilayer Switch Feature Card 2 (MSFC2) is a daughter card on SUP2 that provides layer 3 routing and control functions. MSFC2 has its own CPU and memory to run routing protocols and build the CEF forwarding table, which is then downloaded to the PFC2. MSFC2 is not part of the data path; in other words, it does not perform any per-packet lookups. Per-packet lookups are performed by the PFC2 in hardware.
- The Switch Fabric Module is a crossbar switch and provides two 8-Gbps full-duplex connections to each slot on the chassis for an aggregate switching capacity of 256 Gbps. The Switch Fabric Module is an internally buffered non-blocking crossbar switch. A redundant Switch Fabric Module is supported.

Optical Services Modules

These are high-performance fixed-configuration optical linecards equipped with the Cisco patented Parallel Express Forwarding (PXF) technology to support high-touch line-rate packet services such as destination-sensitive accounting/billing and QoS. The PXF IP Services Processor is an ASIC which contains a 4-by-4 array of CPU cores to perform per-packet service tasks in a parallel and pipelined fashion. The PXF technology offers feature performance as if the features were implemented in hardware, and at the same time provides feature upgrade flexibility via software. The following OSMs will be available:

- 8-port OC-3c / STM-1c POS with 4 GE ports
- 16-port OC-3c / STM-1c POS with 4 GE ports
- 2-port OC-12c / STM-4c POS with 4 GE ports
- 4-port OC-12c / STM-4c POS with 4 GE ports
- 1-port OC-48c / STM-16c POS with 4 GE ports
- 2-port OC-12c / STM-4c ATM with 4 GE ports
- 4-port Gigabit Ethernet WAN

Other Cisco 7600 Components

The Cisco 7600 supports all existing interface modules from the Catalyst 6000 family, such as the FlexWAN module (a module with two slots to support two port adapters from the Cisco 7000 router family), 10/100 Ethernet modules, GE modules, ATM LANE/MPOA modules, and voice modules.

Cisco 7600 Configuration Notes

The following are some basic guidelines on the configuration and operation of the Cisco 7600.

- SUP2 with PFC2 and MSFC2 are mandatory in the Cisco 7600. Redundant SUP2 is optional. The two GE switch ports on a redundant SUP2 can be used just like any GE switch ports on the chassis.
- Crossbar Switch Fabric Module is optional but strongly recommended. The Cisco 7600 has two backplane connections – the crossbar with a 256 Gbps capacity, and the data bus with a 32 Gbps capacity. Without the Switch Fabric Module, there is no crossbar, and all linecards will use the data bus. With the Switch Fabric Module, linecards that are fabric-enabled will use the crossbar, while non-fabric-enabled linecards will continue to use the bus. In other words, a packet going from one fabric-enabled linecard to another fabric-enabled linecard will use the crossbar, while a packet going from a non-fabric-enabled linecard to any linecard, or from any linecard to a non-fabric-enabled linecard, will use the data bus. Currently the only fabric-enabled linecards are all the OSMs and the new 16-port fabric-enabled GE switch modules. Both the WAN ports and the GE switch ports on an OSM use the crossbar. The two GE switch ports on SUP2 also use the crossbar.
- GE switch ports are found on three types of components on the Cisco 7600: OSMs, GE switch modules, and SUP2. GE switch ports on OSMs, SUP2, and the new 16-port fabric-enabled GE switch modules use the crossbar, while GE switch ports on non-fabric-enabled modules use the data bus. Other than which backplane connection they use, all GE switch ports on a Cisco 7600 are the same, regardless of whether they are on the SUP2, on an OSM, or on a GE switch module. GE switch ports are L2 ports, with L3 support via VLANs. All GE switch ports on a Cisco 7600, regardless of slot location, can be used together to form VLANs and Gigabit EtherChannels.
- Each OSM has a PXF-based feature performance capacity of up to 6 Mpps full duplex for high touch packet services such as QoS. The PXF provides packet services only for the WAN ports on the OSMs, including the 4 GE WAN ports on the 4-port GE WAN OSM, but not the four GE switch ports on the OSMs.

- Aside from the availability of PXF services, another major difference between GE WAN ports and GE switch ports is that the GE WAN ports are layer 3 ports and do not support layer 2 functions such as bridging (VLAN) on the GE switch ports. In other words, a group of GE switch ports on the Cisco 7600 can be configured as a VLAN, but GE WAN ports cannot be part of any VLAN. OSM WAN ports (including the GE WAN ports) are controlled and configured by MSFC2, while all GE switch ports (regardless of whether they are on OSMs, GE switch modules, or SUP2) are controlled and configured by SUP2.
- Current OSMs do not perform packet switching; all packet switching is performed by the PFC2.
- Performance of Cisco 7600 depends on which linecards are present. When all linecards are fabric-enabled, and with the presence of the Switch Fabric Module, Cisco 7600 performance capacity is 30 Mpps. When one or more of the linecards are not fabric-enabled, regardless of the presence of the Switch Fabric Module, performance capacity is 15 Mpps. These performance numbers refer to packet switching by SUP2/PFC2. Some of the GE switch modules (not the OSMs listed in this document) support distributed forwarding, and that can increase the overall performance capacity of the Cisco 7600. As a quick reference to gauge the performance requirement of a given configuration, line-rate pps for various interfaces are shown in the following table. Line-rate pps refers to the packet rate of a given packet size that will load up an interface to 100%. So these are "worst case" reference values, as the average utilization of real world networks is far below 100%. From the table, 30 Mpps can saturate 20 GE interfaces with 64-byte packets, or 66 GE interfaces with 256-byte packets.

Table 1

	Line-rate pps (Uni-directional)		
Interface	**64-byte* pkts**	**128-byte* pkts**	**256-byte* pkts**
GE	1,488 Kpps	844 Kpps	453 Kpps
POS OC3c/STM-1c	353 Kpps	160 Kpps	76 Kpps
POS OC12c/STM4c	1,413 Kpps	640 Kpps	306 Kpps
POS OC48c/STM-16c	5,651 Kpps	2,560 Kpps	1,223 Kpps
ATM OC3c/STM-1c	177 Kpps	118 Kpps	59 Kpps
ATM OC12c/STM-4c	706 Kpps	471 Kpps	235 Kpps

* Packet sizes are based on Ethernet encapsulations. For example, a 64-byte Ethernet frame has 46 bytes of IP payload, and when encapsulated in PPP becomes a 53-byte packet when carried over POS; the POS pps numbers in this table reflect that. The same 46 bytes of IP payload becomes two 53-byte ATM cells because of ATM overhead, and that is why the ATM pps numbers in the table are roughly half of the corresponding POS pps numbers with 64-byte packets.

Metro Aggregation with the Cisco 7600 – Aggregation Interfaces

The Cisco 7600 with its GE WAN OSM, POS OSMs, and ATM OSM is ideally suited for access aggregation by service providers in offering metro network services. The 4-port GE WAN OSM is the key building block at service providers' POPs to support the emerging metro GE services, the 8- and 16-port OC-3c / STM-1c POS OSMs provide high-density POS aggregation, and the 2-port OC-12c / STM-4c ATM OSM supports high-performance ATM access aggregation.

All OSMs are equipped with PXF technology to support high-touch line-rate packet services such as destination-sensitive accounting/billing and QoS features, which can be used to provide enhanced service offerings. For example, PXF-based traffic shaping on GE WAN OSM can be used to offer fractional GE services when customers do not have enough traffic to justify a full GE. GE WAN OSM also supports VLAN trunking, so that multiple customers in an office building can share the same GE access switches and access fiber for cost savings, and at the same time achieve traffic segregation and privacy through the use of separate VLANs and VLAN trunking.

Figure 1 shows the use of Cisco 7600 in metro aggregation, terminating various access links including GE, POS, and ATM. On customers' premises, POS and ATM links are usually terminated on routers, while GE links can be terminated on either routers or GE switches. GE links can be easily extended on a campus or office complex through additional GE/FE switches as shown in Figure 1.

Figure 1

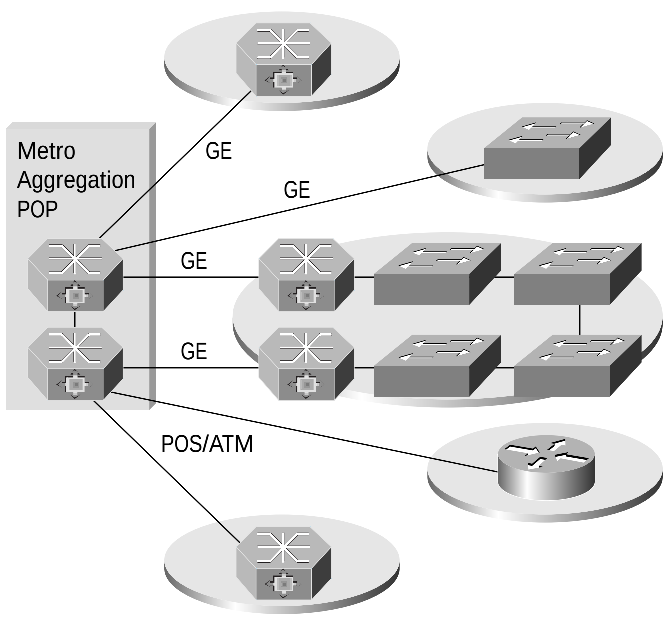

Metro Aggregation with the Cisco 7600 – Intra-Aggregation Connectivity

There is always more than one router in a service provider POP, even within a specific functional layer such as aggregation. There are always multiple routers performing the same function for scalability. The interconnection among routers within a service provider POP, generally known as intra-POP connectivity, is an important subject in any POP architecture design. The goals are performance, scalability, operational efficiency, and modular growth. This section presents an interconnection architecture within the metro aggregation layer and discusses the various design options. The interconnection between the metro aggregation layer and other functional layers within a POP will be addressed in the next section.

The GE switch ports on an OSM are ideally suited for supporting intra-POP connectivity. Figure 2 shows the interconnection of three Cisco 7600 routers using GE switch ports.

Figure 2

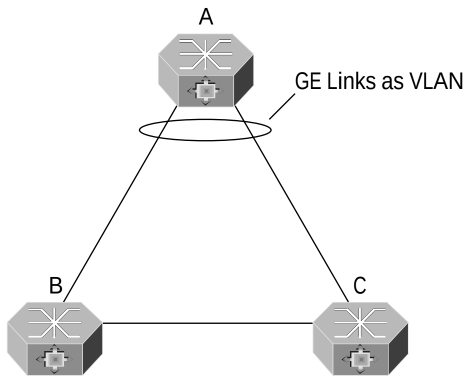

In Figure 2, the three Cisco 7600 routers are interconnected by three GE links. In an ordinary router network, these will be three separate subnets each having its own IP subnet address, thus adding to the consumption of precious IP address space and to overall network complexity and routing protocol overhead. The Cisco 7600 is unique in that it supports both L3 routing/switching and L2 switching. In Figure 2, each Cisco 7600 can have its two GE links configured as a VLAN, so that the three GE links together make up a single subnet with a single IP subnet address. In other words, Figure 2 shows a L2 network of three Cisco 7600 logically connected together on a single GE subnet. The spanning-tree protocol will put two of the three GE links in the forwarding state to form the spanning tree, and the remaining link will be in a blocking state as standby.

It is straightforward to interconnect two or three Cisco 7600 routers as shown in Figure 2. When more units are involved, a systematic interconnection architecture is needed. Figure 3 shows one such architecture.

Figure 3

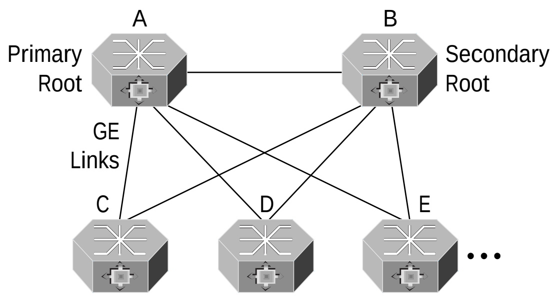

In Figure 3, the first two Cisco 7600 routers, A and B, are configured as primary and secondary roots of the spanning tree. Additional Cisco 7600 routers are then dual-homed to the two roots as shown. In this topology, all Cisco 7600 routers are at most two switch hops away from each other, even if one of the two roots fails. Each active GE link carries traffic to or from at least one of its terminating routers, so no bandwidth is wasted in carrying purely transit traffic. However, because of the spanning-tree protocol, only half of the links (those terminating at the primary root) are active, and the links terminating at the secondary root are not used unless the primary root fails.

It is of course possible to configure each GE link in Figure 3 as a separate subnet, so that the network in Figure 3 becomes a routed L3 network. The result will mean more IP address consumption and more routing protocol overhead, however it does offer the advantage of load balancing on the two physical paths between any two non-root routers. For example, from C to D there are two parallel paths – one through A and the other one through B; while in the L2 interconnection architecture described above, there is only one active path between C and D – the one that goes through the active root router.

A slight variation of Figure 3 can have the advantages of both the L2 solution and the L3 solution. This is shown in Figure 4. Instead of a single VLAN, there are now two separate VLANs, one rooted at A, and the other rooted at B. Traffic from C to D can now be load-balanced (at routing level) between path-A (solid lines in Figure 4) and path-B (dotted lines in Figure 4), and there are no idle links in this configuration. There is only one additional subnet compared to the purely L2 architecture, thus only a minimal increase in IP address consumption and routing protocol overhead, which is a small cost compared to the doubling of interconnection bandwidth. This revised architecture also allows a more efficient interconnection between the aggregation layer and other functional areas within the POP, as will be described in the next section.

Figure 4

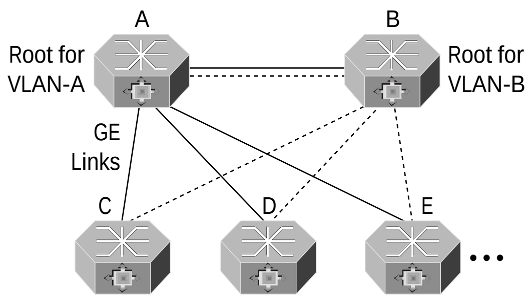

The Cisco 7600 solution offers another significant advantage, which is the support of Gigabit EtherChannels. Each interconnection link between two Cisco 7600 routers in Figures 2, 3, and 4 can be either a single GE link, or a Gigabit EtherChannel which is capable of combining up to eight GE links to form an 8-Gbps connection. Service providers need to support more and more high-bandwidth customers, especially in metro aggregation, and this requires higher-bandwidth intra-POP connectivity. The use of Gigabit EtherChannel is a flexible, high-performance, and low-cost solution to address this intra-POP connectivity need.

Metro Aggregation with the Cisco 7600 – Intra-POP Connectivity

There are three main functional areas within a POP architecture – access aggregation, core backbone (including Internet peering), and local server farms. This paper assumes a segregated POP architecture where each functional area is implemented with a separate set of equipment. So there is an aggregation layer of Cisco 7600s as described so far, a backbone layer of Cisco 12000s, and a server farm layer of routers and switches connected to servers and cache engines. (Smaller POPs can have a consolidated architecture with a single layer of Cisco 7600s performing all these functions, and this is the subject of a separate paper.) The metro aggregation layer has to be connected to the backbone and the local server farms.

The server farms are modeled in this paper as a collection of networks and subnets homed into a pair of gateway routers, which can also be Cisco 7600s. The simplest way to connect the server farm gateway routers to the aggregation layer is to treat the two gateway routers just like aggregation routers, and dual-home them in the same way to the two root routers as shown in Figure 5. Each aggregation router has two parallel paths to each server farm gateway, so there is adequate connectivity for performance and availability. When there are multiple server farms and multiple pairs of server farm gateway routers, each pair can be connected to the two root routers in the same way. Gigabit EtherChannels can also be used to provide high-bandwidth connections between gateway routers and root routers as needed.

Figure 5

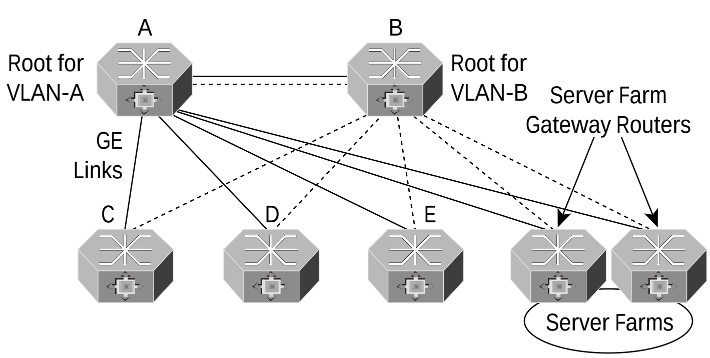

The next interconnection is between the metro aggregation layer and the core backbone layer. The two root routers A and B are naturally the first choice to serve this function. Routers A and B in Figure 5 can be connected to the backbone layer via OC48, Gigabit EtherChannels, or 10GE (future), as shown in Figure 6. Another alternative is to leave VLAN-A and VLAN-B in Figure 5 for just intra-aggregation and server farm connectivity, and use another pair of aggregation routers (such as C and D) with additional GE links to create a new pair of VLANs in a similar fashion as before. Then connect this new pair of root routers to the core backbone layer. The choice of whether or not to add a new pair of VLANs is basically a traffic engineering and scalability issue.

Figure 6

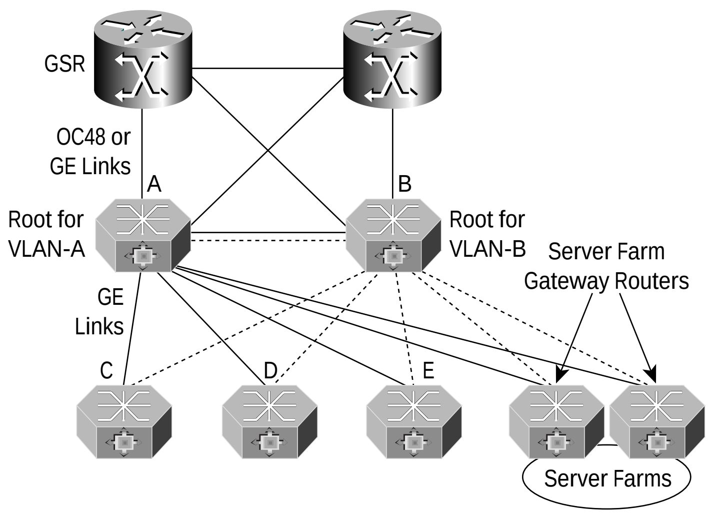

A significant advantage of the architecture in Figure 6 is that intra-POP traffic, which is traffic within the aggregation layer and to/from the server farms, sometimes also known as East-West traffic, is routed exclusively within the aggregation routers. In other words, the intra-POP traffic does not take up any core uplink bandwidth or consume any resources from the core backbone routers. This architecture makes effective use of the unique capabilities and features of the Cisco 7600. It is more efficient, in both cost and performance, than other POP designs where all traffic from an aggregation router goes to the core uplinks, such that all intra-POP traffic has to traverse core backbone routers, thereby increasing the capacity requirements of core backbone routers.

Metro Aggregation with the Cisco 7600 – Traffic Engineering

When there is access aggregation, there is "uplink"; and when there is uplink, there is the issue of over-subscription. Uplink refers to the set of links connecting an aggregation router to the rest of the network within a POP. Over-subscription refers to provisioning a total amount of uplink bandwidth to be less than the total amount of access bandwidth in order to reduce overall network costs. The justification is based on the assumption that access links are not 100% utilized, therefore total uplink bandwidth can be less than total access bandwidth. For example, if there is an aggregate access bandwidth of 10 Gbps with an average utilization of 50%, then 5 Gbps of uplink bandwidth is needed. The 5 Gbps of uplink bandwidth will be fully loaded with traffic from the 10 Gbps of aggregate access bandwidth at 50% average utilization.

A common mistake of over-subscription is double over-subscription. If some of the access links being aggregated are already uplinks from downstream, then those "intermediate uplinks" should not be candidates for any further over-subscription. This is simply because those "intermediate uplinks" will be heavily loaded as they are already over-subscribed with access traffic from downstream.

The first step to size an uplink is to know how much access bandwidth is being aggregated. Table 2 shows the total interface bandwidth for each OSM, counting only the WAN ports (i.e. GE switch ports on the OSM are not counted). For example, a 4-port GE WAN OSM has 4 x 1 Gbps = 4 Gbps. Summing up the per-OSM WAN bandwidth for all OSMs on a Cisco 7600 gives the total amount of access bandwidth that needs to be uplinked. For example, a chassis with three 4-port GE WAN OSMs will have a total of 12 Gbps of access bandwidth. With no over-subscription, a total uplink bandwidth of 12 Gbps is required. With an over-subscription ratio of 2:1, a total uplink bandwidth of 6 Gbps is required, and with an over-subscription ratio of 4:1, a total uplink bandwidth of 3 Gbps is required. To determine the appropriate over-subscription ratio requires understanding of the access traffic load characteristics, and it is important to find out if any of the access circuits are uplinks from downstream. Traffic pattern does change, so monitoring traffic volumes is an important routine in network operations.

Table 2

Optical Services Module	Total WAN Interface BW* (Approx.)
8-port OC3c / STM-1c POS	1.25 Gbps
16-port OC3c / STM-1c POS	2.50 Gbps
2-port OC12c / STM-4c POS	1.25 Gbps
4-port OC12c / STM-4c POS	2.50 Gbps
1-port OC48c / STM-16c POS	2.50 Gbps
2-port OC12c / STM-4c ATM	1.00 Gbps
4-port GE WAN	4.00 Gbps

*For simplicity, bandwidth here is based on nominal interface bandwidth (for example OC-3c/STM-1c is 155 Mbps), and does not account for full duplex, because both access links and uplinks are full duplex. For ATM interface, a 20% adjustment has been included to account for ATM cell overhead.

There are two sets of uplinks in Figure 6. The first is the set of GE links from each aggregation router, such as C, D, or E, to the root routers A and B. The total GE uplink bandwidth from each aggregation router should match the total access bandwidth adjusted by over-subscription ratio. For redundancy considerations, if one of the root routers is down, the GE uplink bandwidth will be reduced by 50%. Acceptable over-subscription during this failed state should be taken into consideration when sizing the total amount of GE uplink bandwidth to each root router.

The GE uplinks in this architecture use the GE switch ports on the OSMs, so one should make sure there is an adequate amount of GE switch ports on each chassis configuration to support GE uplinks. From Table 2, it is obvious that the POS and ATM OSMs have more than enough GE switch ports for uplink, while the 4-port GE WAN OSM has no GE switch ports. The proper combination of OSMs should have enough GE switch ports from the POS and ATM OSMs to cover uplink requirements for the GE WAN ports. For example, a Cisco 7600 configured with five 4-port GE WAN OSMs and four

16-port OC3 POS OSMs has a total access bandwidth of 30 Gbps. Assuming an over-subscription ratio of 2:1, 15 GE uplinks are needed, and the 16 GE switch ports from the four POS OSMs meet this requirement perfectly. If more GE switch ports are needed, 16-port GE switch modules can be added to the Cisco 7600.

The second set of uplinks in Figure 6 is the set of uplinks from the two root routers, A and B, to the core backbone routers. This uplink bandwidth should match the total access bandwidth at the aggregation layer, adjusted by the same over-subscription ratio used in sizing the GE uplinks as described before, and minus the estimated amount of East-West traffic (since East-West traffic will not get to the backbone routers). Same redundancy considerations apply here, as when one of the two core backbone routers is down, the other and the remaining uplinks have to take up the entire aggregation load.

Summary

The Cisco 7600 is an ideal solution for metro aggregation to provide MAN services. This document has presented a reference architecture and design guidelines for metro aggregation, including intra-POP connectivity to other functional layers. The architecture uses the unique capabilities and features of the Cisco 7600 to achieve an efficient, scalable, and high-performance solution. The Cisco 7600 is also an ideal fit in another service provider solution, consolidated service provider POP, which is covered in a separate white paper.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
www.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

DA/0201