Fall 2005

Use of Cross-Connect Clusters to Optimize Routing in Stm-64-Based Sdh Optical Network Systems

Bret William Durrett
Regis University

Recommended Citation

Durrett, Bret William, "Use of Cross-Connect Clusters to Optimize Routing in Stm-64-Based Sdh Optical Network Systems" (2005). *All Regis University Theses.* Paper 356.

Running head: Using Cross-Connect Clusters in SDH Networks

USE OF CROSS-CONNECT CLUSTERS TO OPTIMIZE ROUTING IN STM-64-BASED SDH OPTICAL NETWORK SYSTEMS

by

Bret W. Durrett

A thesis submitted in partial fulfillment of the
requirements for the degree of

Masters of Science – Computer Information
Technology
Regis University

2005

Document Change History

Revision Number	Date of Change	Pages Changed
Draft 2	17 July 2005	Entire Document moved into APA format
Draft 2a	19 July 2005	Changed margins to 1.5” left, Reformatted Certificate of Authorship and Advisor approval pages to fit on a single page each
Draft 3	24 July 2005	Added sections for “Next Steps of the Project” and “What Could Have Been Done Differently” at end of Ch. 3
Draft 4	31 July 2005	Added the Diagram for Cross-Connect Clusters – Pg. 48
Final 5	14 August 2005	Modified Title Page, changed Paper Size to US Letter, Added brief comment to Bibliography regarding first-person project experience
Final 6	22 August 2005	Modified Title Page to remove Clip Art photos
Issued	31 October 2005	Scanned in Required pages with signatures and converted to PDF

ACKNOWLEDGMENTS

The author wishes to acknowledge the Telecommunications Service Provider that I worked for from 2000-2002. The project is based on an existing network within Europe. However, the firm name and any particular characteristics that would identify the firm owning and operating the network have been removed on request of the firm.

Abstract

With the migration to Synchronous Digital Hierarchy, which uses the concept of logical rings, for backbone transmission systems, one of the major concerns that has been brought up repeatedly is a method in which to bring sub-rate circuits from one ring to another without having to decompose the entire backbone data stream to its individual circuits. This is critically important since the backbone data rate can be as high as 10 Gigabits per second or greater and may carry several thousand circuits, ranging in data rate from less than 2.4 Kilobits per second to 2.5 Gbps (STM-16). One potential means of providing this capability in cross-connection locations is to implement cross-connection clusters between the rings. This requires detailed planning of the network infrastructure prior to providing the first customer services, in order to avoid having disruptions to that service at a later date.

As shown in this paper, the consequences for failing to plan and implement a strategy allowing for expansion and flexibility in the network build-out phases can have a significant impact in terms of revenue and reliability later during routine network operations, especially when service is needed for new customers.

Table of Contents

List of Tables

Table of Figures

Glossary

Node. A node is typically a large point of presence where customer circuits are placed on a major backbone fiber. They usually contain large amounts of equipment, redundant power and climatic control systems, are monitored on a 24-hour per day basis for both function and security, and are considered the "nerve center" of a carriers' network. A node is where several high-capacity transmission systems will converge in order to allow the routing of customer data from Customer End Point 'A' to End Point 'B.'

STM-4e. An STM-4e is a 622Mbps circuit that uses coaxial cable and an electrical interface for transmission form point-to-point.

STM-4o. An STM-4o is a 622 Mbps circuit that uses an optical interface and fiber optic cables to transfer data from point-to-point. It is more common to see the fiber optic interface used at this data rate but there are uses for an electrical interface as well.

List of acronyms / Abbreviations

ATM – Asynchronous Transfer Mode

BPS (bps) – Bits Per Second

DSL – Digital Subscriber Line

DWDM – Dense Wave Division Multiplexing

K-, M-, G-, Tbps – Kilo-, Mega-, Giga-, Terabits per second

ISDN – Integrated Digital Services Network

ISP – Internet Service Provider

MTBF – Mean Time Between Failures

MTTR – Mean Time To Repair/Replace

OC-(*x*) – Optical Carrier level (*x*)

POP – Point Of Presence

POTS – Plain Old Telephone Service

PSTN – Public Switched Telephone Network

PVC – Permanent Virtual Circuit

SDH - Synchronous Digital Hierarchy

SLA – Service Level Agreement

SONET – Synchronous Optical NETwork

STM-(*x*) – SONET Transport Module level (*x*)

STM-(*x*)e – SONET Transport Module level (*x*) electrical interface

STM-(*x*)o – SONET Transport Module level (*x*) optical interface

TSP – Telecommunications Service Provider

VPN – Virtual Private Network

Chapter One

Introduction

Technological advances in the areas of telecommunications and fiber optic transmission equipment has fueled an explosive growth in the amount of data that can be transferred in a single fiber optic cable. Previously (prior to 1992), it was considered impressive to carry an OC-12 (622 Mbps) on a single fiber optic cable unless one was a backbone carrier or in an experimental network. It took very specialized (and expensive) equipment, special fiber optic cables, and was reserved for those institutions with large budgets and need for highly reliable and secure data transmission needs. At this time, even major carriers such as AT&T, were operating "high capacity networks" running at STM-16 (2.5 Gbps) and below as experimental systems to prove that the concept would work and be profitable in the business world.

Soon, these systems were brought into operation to carry public data in order to maximize the profits for, and carrying capacity of, the public network infrastructure and the operators thereof. However, as in the quote from the movie "*Field of Dreams*" where Kevin Costner is told, "If you build it, they will come!" the high capacity networks were built and come they did. Users found that they could have data performance at home or in remote locations that rivaled, and some cases, exceeded that experienced when the users were sitting in their own offices. In addition, DSL, ISDN, and Cable modem networks were starting to become available, offering the end users order-of-magnitude increases in bandwidth to the home. In order to cope with the increases in demand, technology found ways to place more and more data on a single piece of fiber optic cable or "glass." Since the buried fiber optic cable represents one of, if not the most significant expenses for the network service provider, be it a telephone company, and alternate service provider or simply a local area Internet Service Provider (ISP), the capability to add capacity to their networks while not having to rip out and replace thousands of miles of cable is an extremely attractive proposition. However, it can lead to problems, especially for those systems that cover large geographical areas or have multiple networks that overlay each other where each set of network media may not “touch” all of the same points of presence. In this case, a transfer of these particular circuits from one network to another is needed in order to get the data from point A to point B. This need to transfer the data from one system to another creates a set of unique issues that were not as prevalent in previous point-to-point network configurations.

Complicating matters further, the implementation of various ways to provide redundant data paths and alternate transmission paths in the event of an outage n a particular segment of a network lead to the formation of a ring structure where the data could be introduced into the ring at a point-of-presence and dropped out again at another POP in a different location. These ring structures were found to offer the most effective methods of redundant paths for data transfer as well as unique “self-healing” capabilities where, in the event of a failure in one segment of the ring, that is to say an outage between any two adjacent points on the network where multiplexing or optical repeaters are located, the data would simple reverse direction, be transmitted through the ring, past its original starting point and to the destination in the opposite direction.

As technology progressed, the composite data rate increased to the point where it is now common to have 10 Gbps or an STM-64 carrier as the base traffic carrier on a backbone network. There can be several of these carriers carried on a single fiber, using a scheme called Dense Wave Division Multiplexing or DWDM. DWDM is simply the use of different wavelengths or “colors” of light to provide different channels on a single strand of fiber optic cable. By adding additional wavelengths of light, the carrying capacity of a single strand of fiber can be increased dramatically. As both fiber optic transmitters and receivers have become more and more sensitive and able to distinguish between finer increments between the wavelengths of the colors of light, equipment manufacturers have been able to take advantage of this and incorporate these new lasers and receivers into their equipment. While the state of the art allows for multiple terabytes of data to be transferred on a single piece of glass, the basic carrier rate has remained at the 10 Gbps or STM-64 level. It is assumed that the 10 Gbps carrier has the most appropriate “bang-for-the-buck” or the best ratio of cost to return on investment, both in terms of scalability and manageability. This makes the 10 Gbps carrier for the backbone service carrier equal to the T- or E-1 carrier in the perspective of the local telephone provider.

Some providers found that it made sense, both technically and for the sake of return on investment, to “layer” several rings logically on top of one another. For example, one ring might be used to pass traffic that was only going to the adjacent POP while another ring could be used to connect the cities with the highest percentages of traffic between them. A third ring could be designated as only available to carry high-rate services (STM-1/OC-3 and above) while a fourth could be designated to be a backup to a special configuration of circuits.

The design of the backbone network is something that is normally done in the initial planning phases. However, even these plans can be a victim of their own success.

In this particular example, the services that were being offered were in such high demand that the preprovisioned network was soon running at full capacity and had to be expanded. This expansion was the driving factor in creating the difficulties. The expansion was done with the simple ideal of getting as many customers online as quickly as possible without considering the repercussions of the needs of those customers. The attitude of “We will make it pretty later, let’s just make it work now.” led to the network having excess unused capacity in some locations and being unable to accommodate customer needs in others. In addition, the nodes in the network were chaotically pieced together in a less-than-optimal fashion, creating additional stress for the technical staff that worked in the nodes.

Basic Assumptions Used

There are some minor but significant differences between the SONET standard used in North American telecommunications systems and the SDH systems used in Europe. One of these is that the lowest level of SDH circuit, which would be referred to as a “Level Zero” circuit (i. e. DS-0) in North America, is the STM-1 or 155 Mbps. For the sake of discussion, I use the convention of calling any data stream smaller than an STM-1 a “subrate circuit.” Conversely, in the North American system, a DS-0 refers to a 64 kbps circuit that has a specific framing and set of control bits included.

In addition, as shown below in Table 1, the names of the specific circuit levels and their associated data rates can lead to some confusion.

North American Circuit Name	European Circuit Name	Approx Data Rate
OC-1	N/A	45 Mbps
OC-3	STM-1	155 Mbps
OC-12	STM-4	622 Mbps
OC-48	STM-16	2.5 Gbps
OC-192	STM-64	10 Gbps

Table 1- U. S. to European Data Rate Conversion

This paper is based on a European network. Therefore, I will use the SDH convention of calling circuits by their STM equivalents.

Issues associated with dispersion of light power, compensated or “doped” fiber optics in this paper are intentionally omitted from this paper since these issues are not relevant to the discussion of multiplexing and demultiplexing in Central Offices or nodes.

Problem Statement

In current ring architectures, it is common that the entire backbone data stream is broken down, the individual circuits that need to be transferred to another ring groomed out, and the remaining circuits recombined and placed onto the next fiber segment for forwarding to the next destination along the ring. This results in excessive amounts of equipment needed as well as increasing the complexity of the wiring inside the local Point-of-Presence. The additional complexity is due to the fact that each individual data stream must be broken down to the circuit level and not all of these circuits are of the same type or bandwidth. For example, a single STM-1 (155 Mbps) could be either an optical or electrical signal. Most often, circuits below STM-1 are electrical (copper wire-based) while those above STM-1 are optical and use fiber optic cable. However, to break down circuits, especially electrical-based circuits, invites introduction of errors into the data stream. These errors can be due to "crosstalk" or the corruption of the electrical signals by external influences, faulty wiring, or most commonly, a faulty configuration or wiring plan. This additional complexity often results in excessive downtime in the event of an outage since the technical staff has to trace the circuit routing through the associated cable trays by pulling the wire hand over hand.

Likewise, this breakdown of circuits requires excessive amounts of equipment, in terms of multiplexers, patch panels and bays, racks, wiring, and by default, climate control and power. Each circuit requires a complete "chain" of equipment in order to be stripped off of the backbone carrier. For those circuits that do not terminate at the node where they are being stripped out, a second complete "chain" of equipment is needed to put the individual circuit back on to the backbone for the transfer to the next node in the system.

One of the issues that quickly became apparent in the new world of ultra-high-speed communications networks was what did one do if one had to take circuits from one ring and move them to another in order to get them to their destination? In any network system, one of the primary goals is to minimize the number of times that a carrier has to be demultiplexed or broken down into its component circuits. The main reason for this is that, for each complete demultiplexing of a carrier, sufficient terminal equipment (lower rate multiplexers or other devices) to accommodate each individual circuit that is carried on the backbone STM-64 must be purchased, installed and maintained. . In addition, each time that a circuit is converted from one format to another, e. g. from Fiber Optic (light) to copper (electrical), errors may be introduced into the data stream. This is a function of the conversion process as well as being based on influence from external sources. The potential external error sources are power, temperature, or humidity. That is one of the main reasons that a node typically has filtered power, shielded lighting, and is kept at a constant temperature and humidity. This implies that, for each circuit “Q” that enters the ring at point A; so long as it does not exit at an adjacent node (point B CW or point (x) CCW), an matching set of devices has to be located at the point where multiple rings converge in order to extract, then remultiplex Q to the backbone carrier rate in order to send it on to the next node on the same ring if another circuit, Q′, which also rides on the same carrier, must be extracted in order to move it to a different ring. See Figure 1 for a simplified diagram that shows this concept.

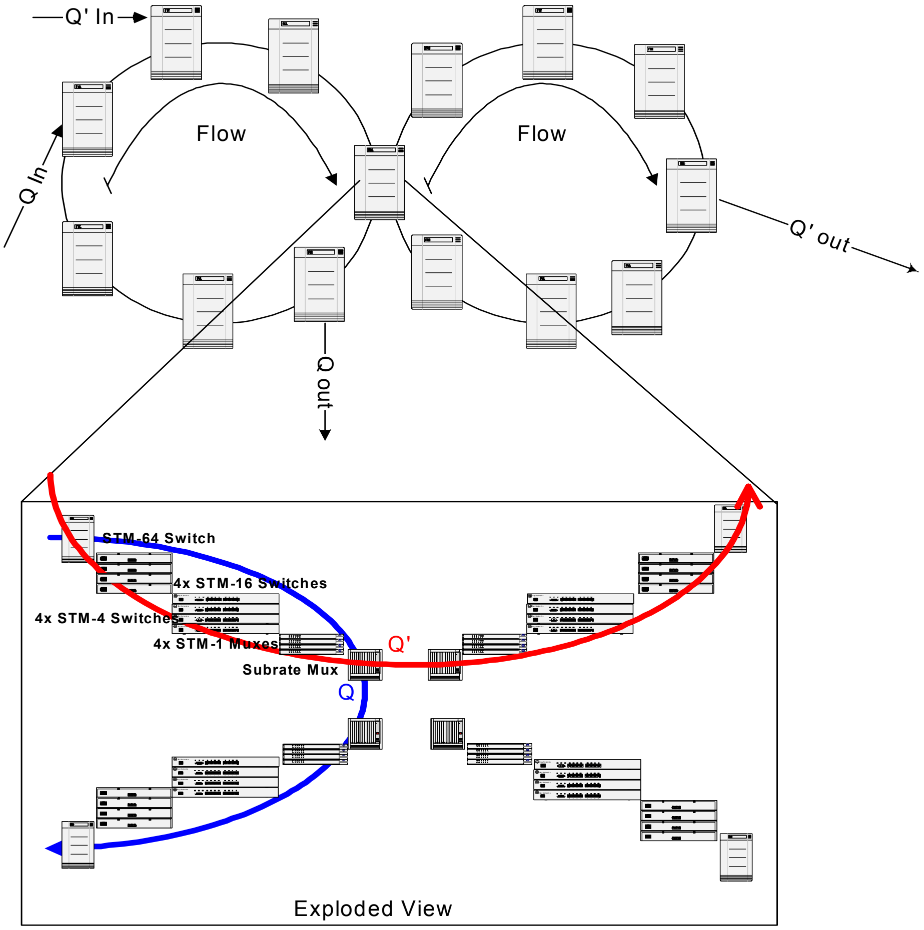

Figure 1 – Simplified Network Breakout Diagram

Figure 1 shows that, when circuits are multiplexed onto an STM-64 (or other high-order carrier) backbone, without regard to the final destination, the entire STM-64 must be broken down into the component circuits in order to extract any single circuit that needs to be transferred to an alternate ring. This requires sufficient equipment to disassemble the entire STM-64 into the lowest level circuits and then rebuild it. This means that, PER RING, there are two STM-64 multiplexers, eight STM-16 multiplexers (four for demultiplexing and an additional four for multiplexing), sixteen STM-4 multiplexers (eight for the demux and eight for the multiplexing), a maximum of 128 STM-1 multiplexers (64 in each direction) and an unknown number of subrate multiplexers used to build the STM-1 circuits. In addition, for each of these multiplexers, patch panels, cross-connection fields, and associated cabling must be present. In a node where several rings are located, the problem be comes readily apparent. The costs associated with the installation, maintenance, and documentation of such a large amount of equipment will have a major impact on whether the carrier is able to provide service at a price that the customer can afford while being able to turn a profit.

In conjunction with the increase in risk associated with the presence of additional equipment, each of the circuit appearances must be appropriately documented. The documentation of a Network Node or Point-Of-Presence is a time-consuming and painstaking task. Documentation must cover each circuit, from backbone trunks at the 10 Gbps (STM-64) level all the way to the individual circuit. Tracking each of these circuits through each of the stages of demultiplexing and remultiplexing, through the various patch panels and cross-connection fields, and through the cable trays associated with the circuit requires precise and extensive documentation. If the circuit is modified in any way, the documentation must be correspondingly changed. In some cases, this could result in changing tens, if not hundreds, of pages. Since documentation is kept electronically as well as in hard copy (paper), ensuring that changes are entered into the documentation system is both critical and time-consuming. In the event that a single change is left out of a single copy of the documentation, the entire documentation trail becomes suspect and is no longer reliable. This creates extreme difficulties for the technical staff at a node in the event of a Move, Add, or Change (MAC) order or in the event of a failure where troubleshooting is required, particularly if there is a conflict between the old and new routing. For example, if a circuit was appearing at Patch Panel A, appearance 5 and, after a change it is moved to Patch Panel W, appearance 198 and the documentation trail for this change is not updated, the potential for confusion and additional problems becomes readily apparent.

There is also a significant cost associated with the maintenance of such documentation. As the amount of documentation increases, so do the costs of maintaining it. While computer-aided systems such as databases and physical plant management systems can help to alleviate some of this burden, there remains a significant amount of work needed to keep these systems up-to-date. In addition to keeping the documentation updated, there is the challenge of keeping the various versions of the documentation synchronized and ensuring that the identical information is placed in all documents related to a specific topic.

In the Telecommunications industry, one typically finds three levels of documentation. These are:

1. Node or POP documentation – These documents describe the physical location, the hardware installed, the cabling present, power and climate control systems, patch panel locations and patch assignments, and information regarding cross-connect or other wiring frame information. This information can be maintained in either hard copy or in specialized computer applications. Often, there will be a so-called “Master” Copy that documents the system as it was built when it was put into operation.

2. Circuit documentation – These documents relate to the individual circuits that transit a site. They may be site-independent and document the entry and exit point of a circuit. They may also be site-specific and provide information regarding the entry of a circuit into a site, the customer information, circuit ID designation, outage information (who to notify if a circuit is reported down), priority of circuit restoral, and other information relevant to a specific circuit. This information may also be part of a paper circuit record or placed in a computerized database.

3. Diagrammatic information – These documents are hybrids of the previous two types but instead of being text-based, are made of diagrams, drawings, or other "pictorial information." Documents of this type would include circuit diagrams, physical plant layout drawings, Rack layout plans and drawings, etc.

In order to keep these three types of documentation consistent, intensive effort is required. By adding the extra equipment needed to break a higher-rate backbone level circuit to its component parts, the amount of documentation required is often multiplied by factors of five or more in order to document the additional equipment, cabling and configuration, depending on how far the backbone circuit is broken down..

Chapter Two

Current Situation

The current situation is that circuits have been placed onto the backbone with very little forethought and a lack of a strategic or long-term plan. Customer growth was unknown as was the distribution of traffic, both in terms of bandwidth and geographic location. In the case described in this paper, the emphasis was to get the customer connection up and running and then to deal with the consequences as an after-thought. This is a typical issue when dealing with networks that were being established by companies that had little or no experience in regional or trans-border networks. In addition, as many of these start-ups were competing with the established (and usually government –owned) monopolists in Europe, the start-ups were forced to offer cut-rate pricing and significantly better service than the incumbents. Because of this, the network which is described here suffered from the following problems:

Inefficient Routing

Some circuits have been routed from their entry point, through the actual cross-over node to the next drop-and-add multiplexer in the ring, broken out there and then put back onto the backbone in the opposite direction, simply because there was no possibility to break the circuit out during the initial transit of the interconnection node. An example of this is shown in Figure 2. This routing is transparent to the customer and produces no degradation of circuit quality but it does add to the complexity of the system, the time needed to isolate and repair a failure, and the amount of equipment needed to carry a single customer (i. e. revenue generating) data stream.

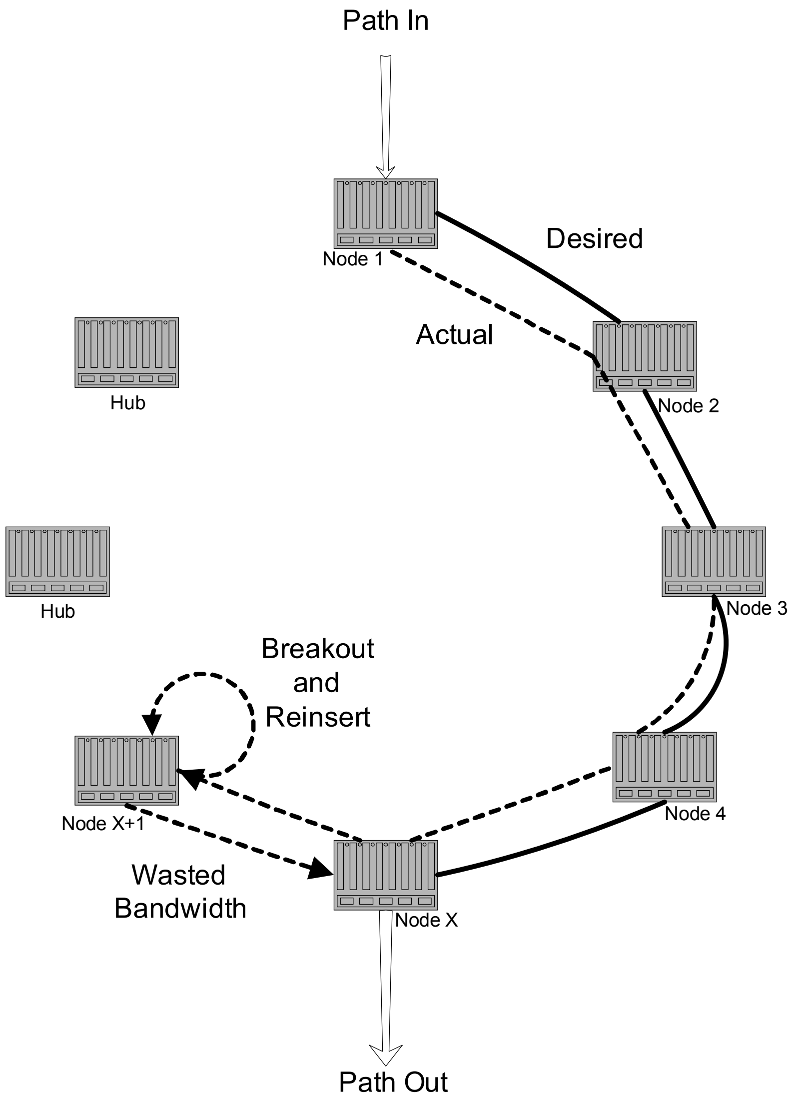

Figure 2 – Inefficient Routing and Wasted Bandwidth

Excess Capacity Usage

As defined in the previous section, the need to pass a circuit through a site and then return it to that same site on a different data stream represents an excess use of capacity. The bandwidth that this data stream is then taking cannot be used to provide service to another customer or to provide additional revenue generation. In this situation, the wasted capacity is a loss to the firm since that capacity is unavailable to sell to another customer and the owner of the data stream itself cannot be penalized for the poor planning decisions made by the carrier by being charged additional per kilometer costs. This is shown as well in Figure 2 as the line marked “Wasted Bandwidth.” This bandwidth, which could have been sold to another customer, must instead be used to return traffic to the appropriate node in order to break it out of the backbone data stream since there was insufficient termination equipment available to break it out in the correct direction. This lack of equipment is strictly due to management decisions that were made without considering all the potential ramifications of such a decision. As such, in terms of immediate expenditures, it is not a catastrophe since, at the time of the decision, the network was not running at capacity or even close to it. Bandwidth constraints were not a major consideration. However, as the network matured and carrying capacity of the network was rapidly being approached, such wasted sources of revenue were made painfully obvious. In some cases, the traffic had to transit more than a single node on the return path, thus exacerbating the problem.

Additional Points of Failure

Not only are the number of potential failures increased, so are the number of locations or points where a failure can occur. Each piece of active equipment in a network represents a potential point of failure. While SDH rings are designed to be self-healing, that is to offer uninterrupted service, the additional complexity and equipment needed to either return a circuit to the correct node for routing to another ring or to break out the entire high-speed backbone link to its constituent parts creates additional points where a failure or outage can occur. These outages can be caused by an equipment failure or by human error. When the amount of equipment and cabling for a specific circuit is increased, the probability of a failure occurring is also increased.

Increased Time To Repair

Finally, the time needed to repair or restore an outage is increased due to the inefficient routing. This is due to the extra time needed to isolate the cause of the failure within systems that would, at first glance, have no bearing on the circuit itself. Because circuits are routed through additional and unneeded equipment simply in order to extract them from the backbone data stream, there are more points of potential failure to isolate prior to determining the actual cause for the outage. These additional and unneeded breakouts were due to the lack of planning that went into the initial build-out and real lack of insight as to the future demands that would be placed on the network. This results in an increase in the time to repair or restore a circuit. The increase in time to repair can be specifically attributed to the extra nodes through which the circuit ran. In one specific instance, an outage at a node that should have been totally unrelated to the specific circuit caused a customer to loose over one hour worth of traffic. This is due to the fact that the engineers at the affected node were simply unaware that their outage affected customer end-circuits. They were operating under the assumption that it was a local level fault which, under normal conditions, would cause an outage of no longer than one second as the traffic automatically re-routed through the node. However, since the customer circuit was broken out there and then remultiplexed into the backbone in the opposite direction, it never reached the redundant switches and was therefore out of service for the duration of the outage.

Financial Impacts

Naturally, the effects of poor planning or poor circuit layout and network design have direct financial consequences for the firm. While this paper does not go into specific details about the profit and losses incurred, the areas in which money was lost or costs increased are identical to other firms operating in the same business. These costs include:

Lost Revenue

Telecommunications companies are in the business of moving data from one place to another. In order to do this, they rent or lease portions of their network to their customers. If a piece of the network is already used to move data due to inefficient routing or poor planning, that part of the network cannot be leased to another customer. Therefore, the company ends up with a potential loss of revenue. If the routing of the data were made more efficient, the additional, unused parts of the network that had been freed from excess traffic could then be leased to a new customer.

Customer Charge-Backs

In many cases, the customers of today's telecom companies are transferring data that is vital to their interests and to their own revenue. In addition, this data can be critical for the firm. For example, one major US brokerage estimated that, in 1994, an outage of the network that they used to trade stocks that lasted for more than five minutes would cost them more than one million US dollars per minute of outage time in lost trading commissions and profits on trades[1]. Therefore, many client companies (consumers of the services) are building charge-back clauses into their service level agreements. These clauses specify the amount that the client can either demand directly from the firm in compensation for an outage or how much the client is relieved from payment on their monthly lease charges. These charge-backs are typically related to the number and duration of outages. Because of the extra equipment and troubleshooting time needed to isolate and restore a failed circuit that is inefficiently routed, it is much more likely that the communications service provider will encounter charge-back situations.

Penalty Charges

Like the charge-back clause, customers are also using penalty clauses in order to have a means of recompensation (or some would say retribution) to recover damages if their circuit is down longer than a specific period. As noted in the previous section, a major brokerage could stand to loose millions of dollars in trading profits and commissions if they are unable to process customer transactions for a space of 5 minutes or more. Penalty clauses in contracts allow the customer to recover a specific amount in the event that an outage lasts longer than the specified period of time. In this case, again, due to the complexity of a mis- or inefficiently routed circuit, the potential that an outage could last sufficiently long to allow invocation of a penalty payment is greatly increased.

---

[1] Study was done by BDM Federal, Inc. Boulder, Colorado for a major US Brokerage House in 1994 as part of a contract to optimize their backbone network. The author of this paper conducted the study.

Customer Service Level Agreement Costs

Telecommunication companies guarantee their customers a specific level of service as provided in the SLA. If this service level, usually based on minutes of service lost in a specific period of time, is not met, the customer has the right to reduce the amount paid for the service by an amount specified in the agreement. If circuits are inefficiently routed or otherwise improperly documented, the MTTR (Mean Time To Repair) will likely exceed that allowed by the contract. Therefore, in the event of an outage that is caused by the company (rather than the customer), the company looses money. If the routing of a circuit is made complex by poor planning, it then takes longer to troubleshoot the circuit, isolate the problem, and repair it. This leaves the telecomm firm suffering the loss of customer revenue. In addition, if the company fails to meet the SLA in multiple occasions (again defined in the terms and conditions of the SLA), the customer may have the right to terminate the service. In this case, the Telco has lost 100% of the revenue stream and will most likely loose additional revenue as the customer begins migrating any other service provided by them to other companies who do not suffer from outage problems. Therefore, it is in the financial best interests in terms of the SLA, for the company to reduce the number of possible failure points to a minimum and to maintain proper documentation as well as to plan the network to route traffic in the most efficient manner.

Additional Cost of Equipment Needed

Each multiplexer in a system is equipped with circuit boards that are designed to perform a specific function. When, due to inefficient routing or capacity mismanagement, additional cards are needed that are not providing a return of cost, either because they are used to “wrap” a data stream back on itself or to break down higher bit-rate carriers to lower levels so that a single bit stream may be extracted, the company must purchase or lease those additional cards. In a large network, that can mean that tens or even hundreds of additional cards are required to provide the contracted service to each customer. These direct costs do not include additional costs such as rental on a physical location to install the equipment since that is covered elsewhere in this paper.

Impact on Equipment

Hardware is required to carry the data from the client premises in one location to the client premises in the second location. This hardware consists of multiplexers, cables, and patch panels. In addition, there are management systems for the multiplexers that are used to monitor and control the performance of the network, test equipment needed to troubleshoot it, and the various hand tools needed simply to put connectors on wires or to remove a screw holding a card into place. If the network is poorly configured, the amount of equipment is usually increased to cover the needs of the network. Since a single circuit could be using more hardware than is strictly required (in comparison to a circuit that is properly planned and efficiently routed), the costs associated with the equipment increase. In addition, one needs a physical location to house and operate this equipment. These "nodes" or locations have specific power, cooling, and access requirements that need to be observed.

Multiplexers

The network on which this project is based was, at the time of the project, consisting entirely of equipment from Nortel, Inc. There were several layers of equipment starting from the backbone Optera Multiplexers (STM-64), to the TN-16 STM-16 second level multiplexers and the TN-4 STM-4 first level multiplexers. The client circuits were normally brought into the TN-4 as a first step, either optically or electrically, depending on the bandwidth and the type of client circuit. From there, the bundled first-level trunk of 622 Mbps was sent straight to the Optera into a low-speed interface or to the TN-16 second level multiplexer where it would be combined with other mid-level circuits to form an STM-16 (2.5 Gpbs). The STM-16 would be sent directly to the Optera DX into a high-speed interface. Each of these interfaces is a single card in the chassis, requiring power, climate control and a "place to live." Each of these components translates directly into costs to the telecom company. Therefore, it is in the best interests of the company to strike a balance between having too much equipment on hand (excess capacity) and not building the network sufficiently to handle potential surges of demand. However, when capacity is used needlessly due to inefficient routing of a circuit, it provides the worst possible combination. Each card in a first level multiplexer can cost over 5000 Euros (approximately $6000.00 at today's exchange rates). This makes it vital that cards and equipment be on-hand ONLY for those circuits that will provide a Return On Investment. For circuits that are poorly routed, this is money that is wasted. If the circuit has to be doubled back on itself due to inadequate capacity in one direction, this can result in the need of a pair of interface cards. Not only was the purchase of two cards unneeded, the costs associated with the cards (power and physical space) is also wasted. Not all circuits that are poorly routed will require a second pair of interfaces to be purchased but the principle remains that there is no ROI for that part of the circuit.

Patch Panels

Between each interface and the "rest of the world," the circuits run through patch panels. These devices are designed to facilitate troubleshooting as well as to allow the ability to "groom" or to place individual circuits into specific ports on the multiplexer interfaces. If additional interfaces are needed, a patch panel will have to be assigned as well to that interface. While patch panels are common and required in telecommunications Central Offices, they can quickly become a confusing array of cables, wires and plugs. Each circuit is patched from input patch appearance to output patch appearance. Circuits that are using excess capacity or are poorly routed at to the confusion because the also require patch appearances. This generates additional work in order to document the appearances of the circuit on the patch panel from where it is extracted from the data stream to the location where it is inserted back into the data stream in the other direction.. Besides creating additional work for the installation team, the additional patches have to be documented as well. Excess patching makes for a documentation trail that is virtually impossible to maintain. Finally, excess patching increases the likelihood of an outage caused by humans. This is because, if there is an undocumented patch cable, it is normal for the technician in the node to remove that patch. If it is documented, it is normally left alone. However, many times, patches are simply left after the circuit for which they were used has been terminated. Therefore, it is not uncommon to see undocumented patches being declared as extraneous and then removed from the patch field. If this particular patch cable happens to be carrying a customer circuit, the circuit is then disrupted.

Cabling Plan

Inefficient routing of customer circuits adds complexity to the cable plant within a node. Additional cables must be installed to accommodate the circuits. This, in turn, requires additional documentation and additional expense, much like extraneous patch panels or hardware. Likewise, it means additional troubleshooting time in the event of an outage in order to trace the routing and locations of the cables in question.

Workarounds in Place

There are several workarounds in place in the network in order to avoid the problems mentioned above. However none of them are fool-proof and each of them has an associated set of limitations. The current sets of workarounds are:

- Bridged Rings,
- Reserved Bandwidth, and
- Ring Overlays.

Bridged Ring Structure

A Bridged Ring is formed when a logical connection connects two physical rings. This connection can occupy any amount of bandwidth up to 100% of the total backbone ring bandwidth. If the connection were to occupy 100% of the bandwidth, a single logical ring would be the result instead of 2 logical rings with a bridge between them. Circuits routed from a location on one ring to a location on the second ring are passed over the bridge in the interconnection location (usually a main node) to the second ring. Figure 3 shows an example of how a Bridged Ring would logically appear.

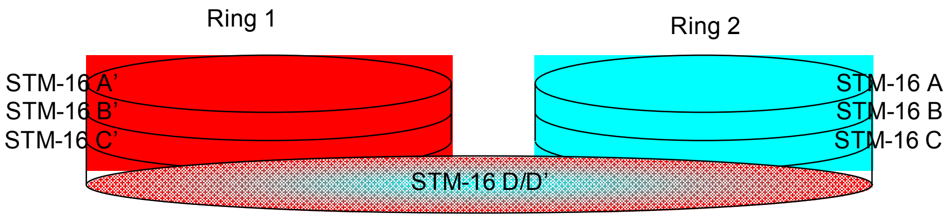

Figure 3 - Bridged Ring Example

Pros and Cons of the Bridged Ring

The Bridged Ring is relatively easy to maintain and document. It is, however, difficult to manage. Because a certain portion of the total ring bandwidth is permanently routed to the “other” physical ring, circuit planning must be done meticulously and the documentation maintained under extraordinarily tight control. In addition, it is difficult to plan the usage of such a bridged system because customer orders are not constant. For example, one customer may lease a 64 kbps circuit while the next may lease an STM-4 (622 Mbps) circuit, both of which may run between the same locations. Therefore, it is often not cost effective to maintain a Bridged Ring system. A fixed amount of bandwidth is allocated according to the capabilities of the multiplexers and, for the capacity that is unused, there is no return on the investment in hardware. At the same time, once the allocation has been made, it is exceedingly difficult to change it without disrupting the other portions of the backbone. This is especially the case of the backbone ring is nearing full capacity or if there are circuits that are logically placed adjacent to the bridge.

One positive aspect of the “Bridged Ring” is that traffic can be routed form one ring to the next at a higher data rate than in some other workaround schemes. Since all traffic in the allocated space is supposed to be transferred from one ring to the next, the allocated bit stream doesn’t need to be demultiplexed to its core components to be transferred. Instead, the entire traffic allocation is passed from one high-speed interface to the next. This scheme does not work efficiently when there are more than two rings to be interconnected due to the need to allocated segments of bandwidth to the bridge.

Reserved Bandwidth

In the Reserved Bandwidth workaround, a specific portion of the ring bandwidth is allocated to inter-ring traffic. In this way it is similar to the Bridged Ring concept but, instead of transferring the entire allocation from Ring “X” to Ring “Y,” the high-speed circuit is broken into its individual low-speed data streams. The individual circuits are then passed to the next ring for remultiplexing into a high-speed data link. In Figure 4, a simplified diagram shows how the Reserved Bandwidth strategy would logically appear. Notice that, for every customer circuit that terminates at a node, a large amount of equipment is needed. In the diagram, the assumption is that all circuits on the particular sets of STM-4’s (both electrical and optical) are to be transferred to the “other” ring (not shown). When this is placed in contrast to making the ring-to-ring transfer on an individual circuit basis, it is clear that this strategy is much better for the TSP in terms of equipment needed, documentation required, and in maintenance to be performed. Likewise, when compared to having no plan (circuit-by-circuit transfer), the Reserved Bandwidth strategy aids in more rapid troubleshooting and circuit restoration.

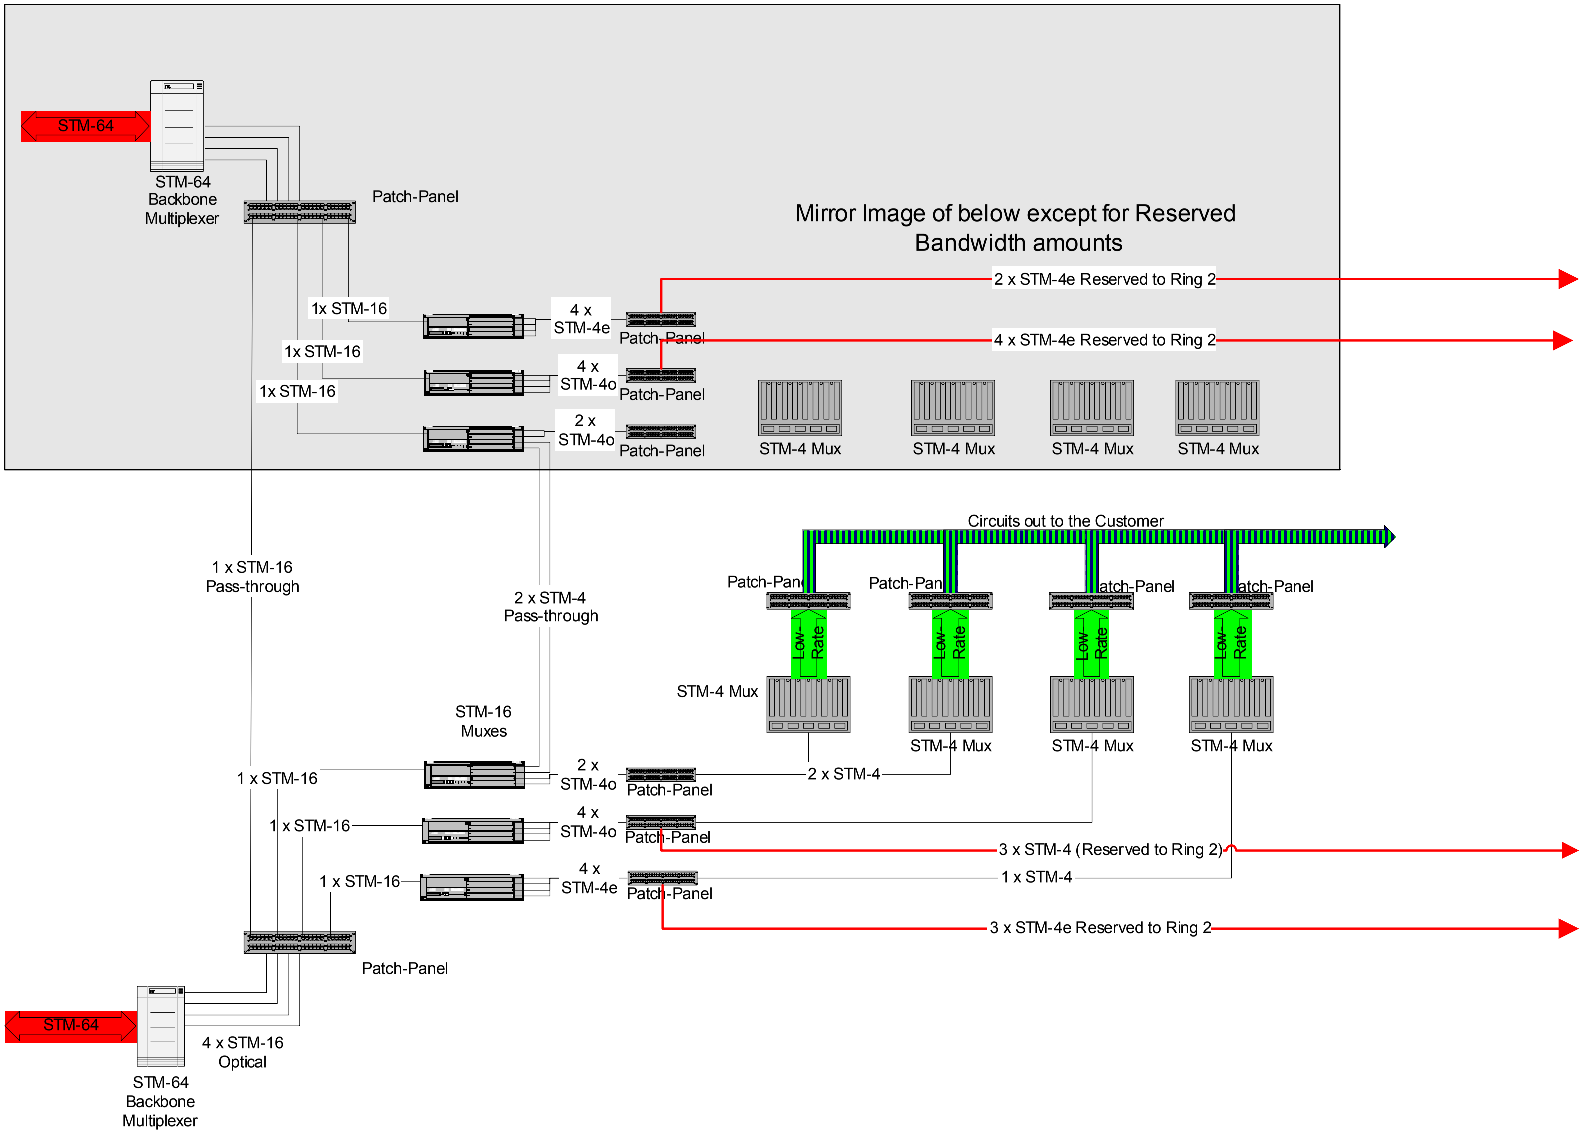

Figure 4 – Reserved Bandwidth Example

Pros and Cons of Reserved Bandwidth

Under the Reserved Bandwidth scheme, more than one interconnecting ring can be serviced within a single node. This provides the carrier with more flexibility in their circuit design, traffic planning, and allocations. However, it carries with it a high price. For each of the constituent circuits that are demultiplexed, a corresponding low-speed interface port has to be provided on BOTH rings (the exiting as well as the entering ring). In addition, cable and patch panel appearances must be provided. Since the number of circuits in a high-speed data stream of this type could easily be several hundred, the scope of the equipment, cabling, and documentation becomes apparent. In addition, by reserving a portion of the backbone bandwidth specifically for inter-ring traffic, the risk of underutilization is present. Since traffic that is bound for locations further along the same ring cannot be routed within the reserved bandwidth (as long as the reservation is respected), there is a further risk of both underutilization of the reserved space and insufficient capacity on the remaining bandwidth.

Ring Overlays

Ring Overlays are a special case and are built by adding an additional ring on top of an already existing ring. Adding an additional STM-64 backbone that is routed identically to one previously installed gives an enormous amount of flexibility but requires additional equipment with the attendant drawbacks (cost, maintenance, documentation) as well as requiring customer circuits to be moved (meaning downtime, although it may be only momentarily). In an overlay, only those circuits that will transit to another ring are routed. This means that, for each city attached to the ring, there can be a specific amount of bandwidth reserved and allocated. At the central node where the multiple rings meet, the entire backbone is broken down to its constituent components and each circuit is routed to the appropriate ring for transmission to its final destination or to the next major node. A Ring Overlay system is used to add capacity to an already existing ring structure. If looked at in other terms, a Ring Overlay would look like an additional row of bricks on an existing wall. An example of what a ring overlay might look like is shown in Figure 5.

Pros and Cons of Ring Overlays

Using a Ring Overlay approach allows the network provider an unparalleled degree of flexibility, simply due to the fact that additional capacity is available. However, this capacity can soon be subject to the same problems that the initial ring suffered if not closely maintained and properly planned. Depending on the number of cities attached to the ring and the number of circuits being moved onto it, this solution can offer a cost-effective means of providing inter-ring transport services. The number of circuits being routed directly influences the Return-On-Investment. For those cases where there are many circuits or the circuits require higher bandwidth, the ROI can be substantial.

The opposite can also be true. Since a complete new fiber is being lighted, the company is faced with the purchase of a complete set of equipment, including repeaters and regenerators, or, at the least, enough backbone-level multiplexers to provide the input to a DWDM fiber optic system. Considering that the typical Nortel Optera DX (an STM-64 multiplexer commonly used in Europe) costs well over €100,000 when equipped with redundant power supplies and a full suite of both front and backplane circuit boards, and that each city or connection to the ring will need two of these devices, the costs can quickly achieve levels which are no longer commercially viable in terms of Return on Investment. In cases where the inter-ring traffic is low, the investment in equipment cannot be recovered during the lifespan of the equipment, even when considering capitol depreciation and income from leasing of circuits. In addition to the cost of the additional equipment, the ROI of any currently installed equipment is also impacted. The bandwidth once being charged and accounted to Equipment Set 1 is now accounted for differently. Unless additional local (non-inter-ring) traffic can be found that compensates for the loss of income provided by the inter-ring traffic that is moved, the situation ends up being that both rings are underutilized and are a drag on profits for the firm.

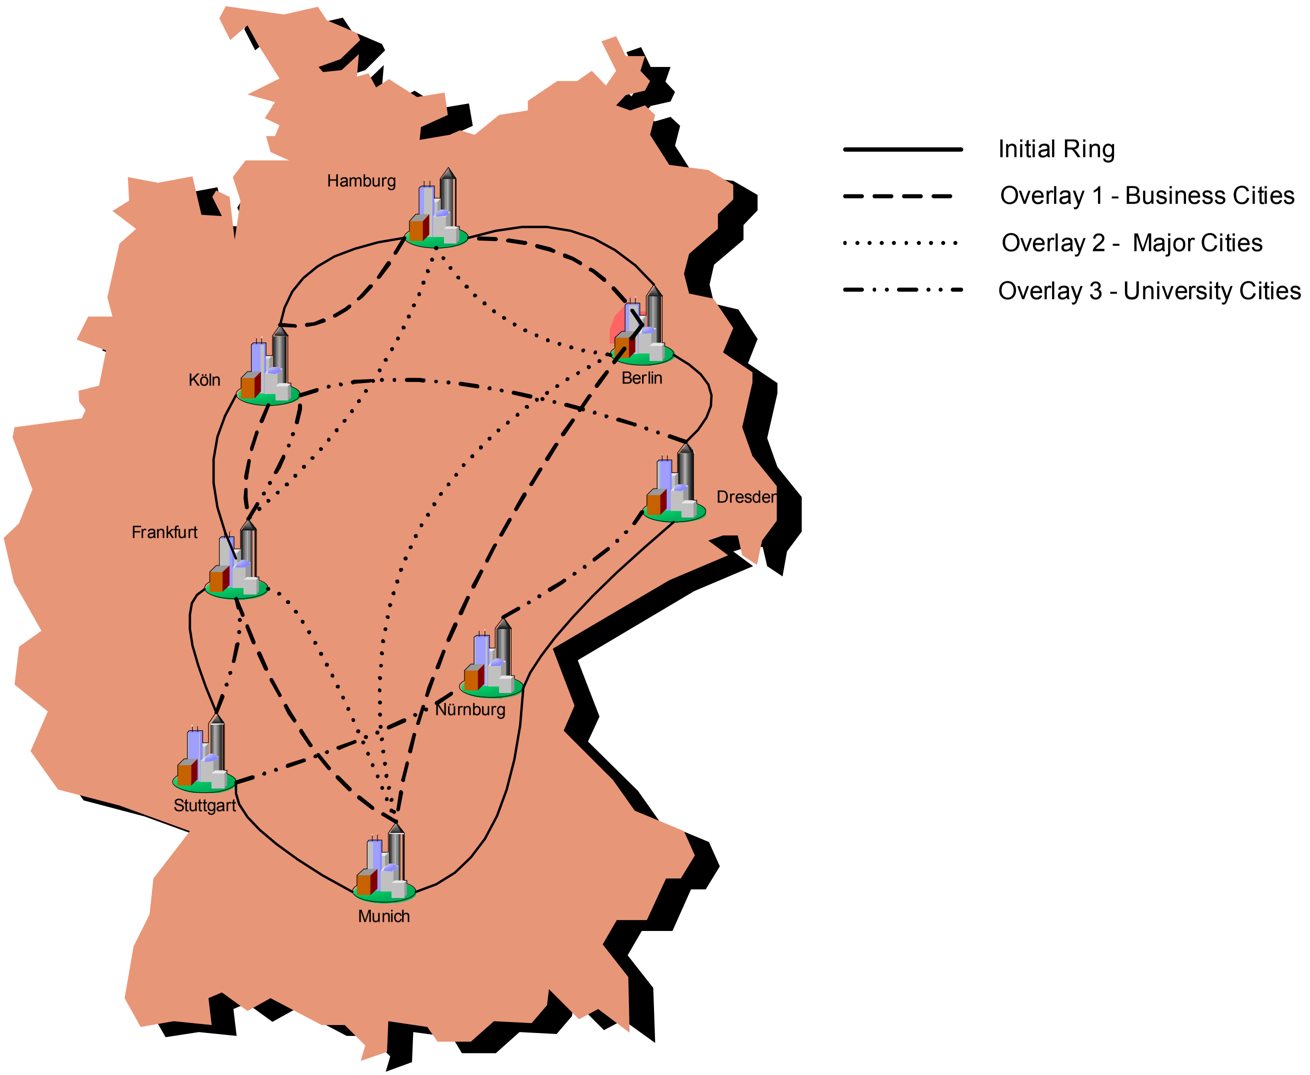

Figure 5 – Ring Overlay Example

Proposed Solution – Cross-Connect Clusters

The solution that was chosen was to implement a ring overlay scheme as defined above but to equip the nodes where the rings congregated with what was called a "Cross-Connect Cluster" which operated at an STM-16 (2.5 Gbps) rate. The Cross-Connect Cluster was a miniature ring structure in its own right that had interfaces to each of the rings in question. The benefits of the Cross-Connect Cluster were that low-speed circuits did not need to be broken down out of higher bandwidth second or third level backbone trunks (STM-1 or STM-4) in order to transit from Ring 1 to Ring 2. The entire STM-16 could be extracted from the STM-64 Backbone, routed optically into the Cross-Connect Multiplexer (a Nortel TN-16) and then the 2nd and 3rd order trunks would be routed to their respective ring interface. This functioned as a combination of the reserved bandwidth scheme combined with the ring overlay scheme and is shown in Figure 6.

Pros and Cons of Cross-Connect Clusters

The Cross-Connect Cluster system allowed any number of rings to be interconnected as long as the composite bandwidth of all traffic on the Cluster did not exceed STM-16 rate. In the cases where more than a single STM-16 was transiting from one ring to another, the transition was made optically between the Optera DX's. The Cross Connect Clusters were specifically designed to accommodate the enormous number of lower-rate circuits that would be provisioned on to a trunk.

In addition, not only did the Cross-Connect Cluster allow transitions between rings serving different regions of the world (which was the initial purpose), it was also found that traffic could be routed from one ring to another within the overlay in nodes where traffic was loading a single ring segment disproportionately to the rest. This served to alleviate congestion in certain segments of the ring where traffic would be simply passing through a node rather than terminating there. By using the Cross Connect Cluster concept, the carrier was able to perform a “load-balancing” between rings that served a single country as well as being able to efficiently route traffic from rings serving one area to rings serving another.

Another benefit of using the Cross-Connect Clusters was that the amount of cabling and patch appearances required was drastically reduced. By running a composite STM-16 from the backbone multiplexer to the input to the Cluster, a single set of patch appearances at an optical level were required. All other routing took place internal to the Cluster equipment. This means that, for each backbone ring that was connected to the Cluster framework, a single pair of patch panel appearances were required. In addition, some equipment, such as patch panels, could be eliminated by taking the direct output of a multiplexer (i. e. an STM-4 “circuit-side” output of an STM-16 multiplexer) directly to the input (trunk side) of a lower level STM-4 multiplexer. The low-level multiplexer was then used to further break the circuits down to the customer level. Contrast that to the number of optical and electrical appearances needed to break down an entire STM-64 backbone segment to its component parts in order to extract a certain number of low-rate circuits and the benefits of this approach become clear, both in terms of effort required and cost. Furthermore, the documentation requirements are drastically simplified for the same reasons as previously mentioned. Every set of patch appearances or cross-connects (on an electrical level) than can be eliminated reduces the burden of paperwork for the TSP.

The verification and validation of the concept also permitted the carrier to accept traffic from cities that were not an integral part of their own network. In several instances, a city would be contractually bound to a specific local carrier. This contract gave exclusive “last-mile” rights within the city to the carrier. These contracts did not extend to long-haul services but, due to the very nature of the contracts, long-haul carriers were prohibited from providing direct connections to the customers in those cities. The Cross Connect Clusters fostered the development of the Point-Of-Presence City or PoP City concept where, the local carrier would provide an interface at the outer boundary of their contractually exclusive area via SDH Optical carrier. They would then route all the client data to that interface from the various locations through the city where we would pick up the data and then transport it to its appropriate destination.

In order to provide this transport service without installing a substantial amount of equipment at this network interface, a drop-and-add multiplexer would be installed on a subrate (STM-16 or lower) trunk. These trunks would be routed to the nearest Node on our network. We would then multiplex the data into our backbone and route it through our network. Since the subrate trunks were usually less than 100% full, it didn’t make sense to bundle the entire trunk into a high-rate backbone stream, especially since there were multiple customers that used a single subrate (STM-1 or STM-4) channel and their data didn’t all terminate in the same locations. Therefore, a cross connect cluster would be implemented at the neighboring node. The data would come in from the PoP City trunk, be “routed” to the appropriate ring via the cross connect cluster, then multiplexed onto the high rate backbone. The same principle, although in reverse, is used to send data to the PoP City customers.

Naturally, there were drawbacks to this approach as well. These drawbacks were far outweighed by the benefits realized. However, the drawbacks were that:

1. Additional equipment had to be installed in the existing nodes. This means that additional space had to be found within the existing rack planning in order to accommodate the new equipment. However, when compared to the amount of equipment needed to implement some of the other solutions, this was found to be a lesser consideration.

2. Additional documentation had to be prepared. For each circuit that transited a cross-connect cluster, the appropriate circuit documentation needed to be prepared and maintained. Again, this was found to be a minor consideration because the majority of the switching was done at higher rate (STM-1 or STM-4) level instead of single circuit level. This had the end effect of reducing, rather than increasing the amount of documentation needed for a particular node.

3. The workload on the Network Planners increased temporarily. The Network Planning group had to plan for the cut-over of all the affected circuits, including planning the new routing to the level of optical port assignments, cut-over times (which included scheduling downtime with the customer), and ensuring that the responsible people were actually on-site at the time of the cut-over. Since most customers needed their data to be available during working hours, the majority of the cut-overs occurred then in the night.

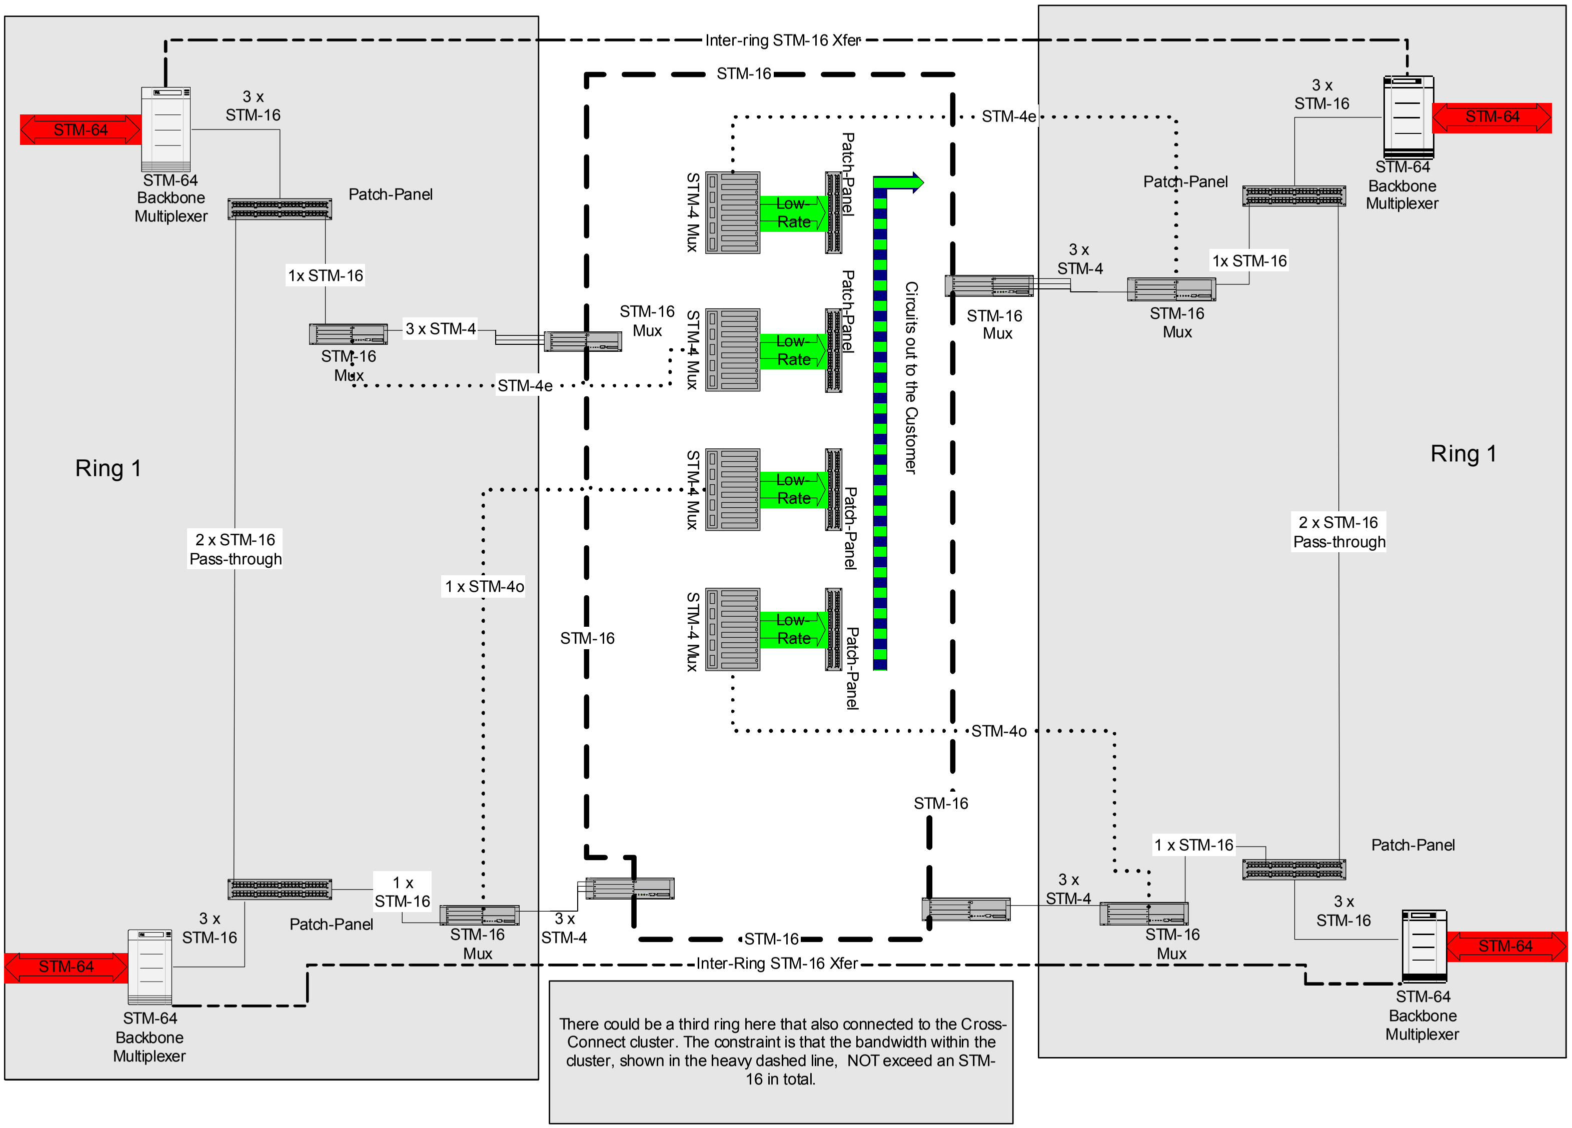

There could be a third ring here that also connected to the Cross-Connect cluster. The constraint is that the bandwidth within the cluster, shown in the heavy dashed line, NOT exceed an STM-16 in total.

Figure 6 – Cross Connect Cluster Example

Chapter Three

Conclusions and Lessons Learned

Conclusions

Several conclusions can be drawn from the issues discussed above. First and foremost, in my opinion, is that, in order for a network to be efficiently operated and to maximize profit for the network owner/operator, the majority of the planning for the network design must be prepared long before the first equipment installations are made. After seeing the effects of the various implementations of the schemes described, it is apparent that some planning prior to critical installations would have reduced the effects and numbers of “reactionary” results (opposed to proactive) that resulted in wasted time, wasted money, and loss of customer goodwill.

The second conclusion to be drawn is that there are more considerations in the development of an overall backbone network structure than were first taken into account. These considerations include market forces, competition, regulation, and customer requirements. In the case of the network described in this paper, shortcuts were taken in order to get the network up and running in the shortest time possible. These shortcuts were felt by management to be in the best interests of the firm in order to begin showing revenue as a return on the investment of capital equipment and installation costs of both equipment and fiber optic capacity to the shareholders and investors.

These shortcuts were found, in reality, to have caused increased downtime for customers in the event of a failure, increased costs over the long term by requiring additional personnel, increased costs in terms of additional documentation production and maintenance, increased costs in terms of additional equipment, both for initial purchase and for maintenance, and to have been a prime factor causing the overall chaos that was seen in virtually all major network nodes.

The third conclusion that can be drawn is that, on a purely technical basis, the cross-connect cluster concept best serves the needs of a growing and dynamic network interchange structure. Networks that are small, static, or that have a fixed set of data communications needs can use one or a combination of the other methods described as workarounds in order to efficiently route traffic from one place to another but in the changing environment of a major network carrier where data traffic is in a constant state of flux, the additional flexibility of the cross-connect clusters allows the most efficient routing of network traffic across various rings while reducing the amount of extra equipment and documentation required.

Finally, a fourth, and somewhat off-topic, conclusion that could be drawn from the various problems and solutions presented here is that a robust Network Planning Tool that allowed simulation of various schemes, if implemented early in the planning process, could significantly reduce the amount of "trial-and-error" needed in order to optimally configure a network. This tool would ideally allow simulation of the different types of connection/interconnection capabilities provided by the equipment vendor of choice but would not be locked to a specific vendor. This would allow the Network Provider to simulate the impact of introducing a new piece of equipment, a new node design, a new interconnection scheme, or potential consequences of failure prior to making a large capital investment in equipment.

Lessons Learned

Since the initial build-out of the network and the subsequent expansion into a total of 32 European cities, cross-connect clusters have been implemented in all nodes where more than one backbone ring has a drop-add presence. Rings in an overlay structure have their own cross-connect clusters, especially in those cities where one or more rings may be nearing full capacity. In addition, when new node was being installed, a set of cross-connect clusters was installed as part of the initial build-out in order to facilitate the routing of data from ring to ring at a later time.

One of the main lessons learned is that it is easier, less expensive over time, and much less disruptive to the network to install the additional equipment required for a cross-connect cluster during the initial equipment installation than it is to retrofit a system in to an existing node. Since the benefits far outweigh the drawbacks of having such a system, a management decision was taken to add the cross-connect cluster(s) to the initial build-out plan for any new node where more than one ring was present and in which those rings had a drop-add point.

Additionally, nodes that were already in operation would also implement cross-connect clusters when practical with particular emphasis on the larger nodes in the main cities where potentially five to ten rings could transit, each with their own drop/add points. These clusters would be installed and tested prior to having customer traffic rerouted onto them. This rerouting was, for the most part, transparent to the customer since it took place at night, during weekends, and/or over a period of several weeks in order to have the least possible impact on customers' service. Since this decision was made after the nodes were operational, the migration of service onto the new system was a slow and pains-taking process.

The most important lesson to be gleaned from this is that substituting flexibility and quick reaction times for thorough planning in the beginning stages of the project will often result in significant problems in the later stages. Once the hardware, cabling, and documentation is in place, regardless if there is customer traffic on the system or not, it becomes much more difficult to modify the system. Add the requirements levied by the carrying of customer data to this and the problem quickly takes on epic proportions. One of the major reasons for the difficulties is that managing the customer can be very cumbersome and difficult. When the customer requires their services to be available on a constant (24x7) basis, the smallest interruption or request for interruption can result in an enormous amount of work. Some of this work would be to provide an alternative path for the services, the customer may choose to switch to a back-up system that they have in place, they may request that such an interruption be limited to a specific time and date and limit the amount of time the service can be down, etc.

A good deal of this reticence on the part of the customer can be explained by the fact that, to them, the data is an integral part of their overall business and business strategy. Customers such as news organizations, stock markets, banks, and brokerage firms rely on their data getting to its appropriate destination as quickly as possible in order to make a profit. When confronted with an outage, planned or unplanned, these firms see a direct correlation between the loss of connectivity and a loss of profit. For these firms, while a cutover to new system may bring them better performance, lower latency, and an increase in reliability in the future, in the short-term, they are confronted with an outage.

On the opposing side of the issue, the Service Provider wants to maximize profits by providing service to these customers. However, when faced with a maintenance or repair issue, the Telecommunications Service Provider wants to resolve the problem as quickly as possible in order to avoid a catastrophic failure on (or in) a node, individual customer needs notwithstanding. The Service Provider is concerned with the overall performance of the network and how actions will impact the overall user community. In this case, one could say that the providers have a “Global” perspective as opposed to the “local” or “self-centric” version of the end customer. From this point of view, repair of the problem or installing a work-around takes priority over individual customer services.

Balancing these competing needs is required when dealing with a system that is already on line and active. This is the prime reason that network design prior to installation and activation is so important. By actively engaging in planning prior to installation of equipment, the resulting network infrastructure can be made robust and have sufficient flexibility. In this particular instance, this planning was not done prior to implementation and the result was that it took a greater expenditure in terms of time and equipment than was strictly needed.

Next Steps of the Project

Writing this section is, for the author, impossible due to the fact that the author was laid off from this firm as the retrofit of the nodes began. In the first wave, over 10% of the total head count was cut. In 2002, the firm had over 5000 employees throughout Europe. Since that time, there have been additional lay-offs and the head-count has been lowered to approximately 3800. The company is currently off-shoring a great deal of the help desk and planning functions to India in order to further reduce network operating costs. In the view of some employees, this is the prelude to the firm being sold or merged.

However, having kept in contact with some of the people who worked with the author during his employment, the author has learned that the cross-connection structure was implemented in nearly all nodes across the network. One of the main reasons for this was that in each city where a company network was built, a node was also built. The city networks (also STM-16 or STM-64) were also SDH rings and were therefore, simply "miniatures" of the backbone ring structure. In these cases, the cross-connect cluster system also made sense for the exact reasons it was appropriate for the long-haul, or backbone, network.

The city networks were SDH networks that were established in all cities in which the company had customers outside of those served exclusively (legally and contractually) by a firm other than the incumbent Public Telephone System. These mini-networks functioned identically to the backbone system in that the customer traffic was collected and multiplexed onto higher speed trunks. These trunks were routed throughout the city and customer traffic that was only point-to-point within a city would then be broken out at the city node nearest the customers' end-point. Those circuits that were destined for other cities or, indeed, other countries were sent to the backbone node within a city to be placed on the wide-area backbone. The city nodes were under the same constraints in terms of size and capacity as the backbone nodes although they had not encountered the difficulties to break out circuits in the appropriate node as had occurred in the backbone. For the cities, implementing a cross-connect cluster at the city network-backbone network interface was relatively simple. For them, it was a simple matter of circuit "grooming" that was able to be done on a logical basis rather than a physical patch or other service-disrupting measure. A logical move of a circuit happens electrically/electronically and is commanded from within the network management system. The result is that, instead of leaving a multiplexer on card (x), port (y), it is routed to card (q) port (z). For the customer, this is totally transparent, assuming that the port has been correctly preconfigured to enter into the cross-connect cluster.

By expanding the cross-connect cluster strategy to the local city network interfaces as well as to the intra- and inter-country backbones, the company changed the most basic dynamics of customer service provision. Instead of having a MAC order that required several stages of documentation review, cable tracing, patch panel re-routing, and down-time associated with it, actions were taken by the local network management system operator and the circuit rerouted electronically after any pre-provisioning work had been done. For the customer, this was ideal since they were either not impacted at all or saw the switch as a momentary "hit" on the line that caused no data loss. A comparison to a "line hit" in the POTS would be hearing a "click" on the line while one is talking to a friend in the next city.

In addition, the company expanded into other cities within Germany that had exclusive arrangements with city-owned TSPs by creating a Point-of-Presence, or PoP node. This was a simple add/drop location that was a stub network off of the nearest major node. Traffic would be delivered to the PoP node and then handed off as individual circuits to the city-owned TSP. Likewise, traffic was delivered from the City-owned TSP to the PoP node, multiplexed, and then transferred to the next node.

The cross-connect cluster strategy served to allow the company to continue its flexible and responsive approach to the customers, quickly allowing customers to order additional services, reroute existing services, and to increase traffic demands. It also allowed the company to increase the scope of their offerings. From relatively simple data and voice offerings, additional services such as point-to-point LAN, wide-area LAN (although this is a complete contradiction in terms), and managed network services are now being offered. When consideration is given to the fact that the majority of the competitive Telecommunication Industry in Europe is truly in its infancy, the ability of a private and independent TSP to be able to break into the business is quite impressive. One has to take into consideration that, at the time of the AT&T divestiture in the United States, virtually ALL telecommunications Services in Europe were still being provided by government-owned and controlled entities. In Germany, the incumbent was the state-owned monopoly, the “Deutsche Bundespost.” It was not until the late 1980’s and early 1990’s that the European Telecommunications Industry was allowed to migrate to a more open and competitive environment. The fall-out from that migration is that several companies, not the least of which were major global players such as Level 3, Global Crossing, MetroMedia, and WorldCom, entered the European market in a flurry. What then transpired was a virtual free-for-all that left many competitors battered and bleeding financial red ink by the road side. The barriers to entry were very severe and resulted in a huge amount of work for the local regulatory authorities as the incumbents, who were a) forced to compete for the first time, and b) were loathe to see their monopolistic profits vanish, set nearly criminally high prices for use of the “last mile” (the part of their existing network that connected the individual customer to the nearest switching station), onerous co-location agreements and space rental costs that were simply designed to prevent a company from being able to exist. After loosing a series of regulatory judgments in Germany, for example, the Deutsche Bundespost was broken into two firms and privatized. One of these firms, Deutsche Telekom, remains the incumbent telephone company, although this is also beginning to change as other firms such as Arcor are providing POTS and DSL services to the individual home user.

The firm which runs the network that is the subject of this case study is one of the few (there are fewer than 10 throughout Europe that are still surviving or have not withdrawn to their own home markets) new-comers into the market that are present in nearly all the European Union countries, have networks that transverse borders of countries without the regulatory and bureaucratic nightmares that are present with the monopolists, have had sufficient cash reserves and return on investment from current customers to weather the cut-throat competitive environment fostered by the European Communities' decision to allow full and open competition, and have been able to circumvent the difficulties placed in the way by the incumbent monopolies, largely by having built their own fiber optic infrastructure, nodes, and networks.

The firm was able to achieve this position by being flexible and responsive to customer needs, providing a level of service that the incumbents could not for a price that was far lower than that of the incumbents, providing services that the incumbents could or would not, and, as a specific selling point, was one of very few companies that was able to offer a complete network solution out of a single firm. The city networks allowed the company to literally run fiber to the basement of the building in which the customer resided so there was no dealing with additional service providers or the local Telephone company. It was the beginning of the “one-stop shop” for Telecommunications Services, which, to that point, was a virtual unknown in Europe.

The future of the firm is unknown. Rumors within the company have it being sold off by the investors anywhere from the end of 2005 to the end of 2006. The major investor naturally can not and will not comment on any potential dealings in order to not affect the price of the company stock. Financially, the company expects to achieve their first cash-flow positive quarter in the 3rd quarter of 2005. This has been the goal of the company (to be come cash-flow positive in Q3 2005), since the year 2000 when the author joined the firm. If this goal is, in fact, realized, it will mark one of the first successful entries into the TSP market in Europe of a firm that was completely independent of the incumbent monopolists.

What Could Have Been Done Differently

One of the author’s major causes of stress and distress during his time of employment and working on this project was the apparent lack of planning that was evident in virtually all stages of the network implementation. There were several occasions where the author was called on to deliver a report of the current capacity of the network only to have management say that the numbers were wrong. The reports were based on the information that was available to the author at the time from various sources. The problem was that there were different sources and virtually none had the complete picture. One source was the circuit database that was maintained at the Network Management Center. This database was used by the engineers provisioning the network and who were responsible for actually assigning ports and cards in varoius systems to transport the customers’ data from node to node. Another source was the network planning database that was supposed to show the loading of the overall backbone on a segment by segment basis. The third source was the network management system itself that showed what was active and in service at any one time. The major problem was that none of these sources agreed with each other and none of them contained the actual status of the network as used.

In one specific instance, the author was tasked to route a pair of STM-4 circuits from Copenhagen, Denmark with one terminating in Frankfurt and the other terminating in Brussels. Based on the information from the circuit database, the planning was made to go from Copenhagen to Hamburg, transition the Hamburg node and then split the two circuits in Frankfurt with one going to the city network and the other crossing to the European Network to be terminated in Brussels. While this would seem to be a straight forward routing, the author then proceeded to spend the next 12 hours dealing with the various locations calling him and saying that the port or card specified in the routing request was either in use, non-existent, or not of a type compatible with that particular circuit type. This was, unfortunately, NOT an isolated incident.

While not the subject of this case study, the author strongly believes that the vast majority of these issues could have been avoided if a proper network planning system had been put into place. One of the authors' tasks prior to being laid off was to research and recommend the implementation of exactly such a system. One system was found that would interface with both the network management system directly and with the circuit order database to show the real-time status of the network at any time. In addition, the network planners could designate cards and ports that were reserved for projects, assign virtual routing of circuits throughout the network to ensure that there was sufficient capacity in the network, both on the primary traffic carrier and on the backup fiber, and to provide usage information on a network, ring, or segment basis. In addition, the system that was recommended by the author was independent of the type and brand of equipment that was in place in the network. The system had logical interfaces to virtually all the major network equipment manufacturer's management systems without modification.

Unfortunately, the authors' direct manager found the costs of such a system too expensive and showed a decided preference for a system that would be provided free-of-charge but was restricted to a single manufacturer of network equipment. This system was researched and found to be totally inadequate for the tasks outlined but, as managers are sometimes wont to do, the manager stayed with his preference. In the end, no system was ever fielded. This caused an enormous amount of difficulty for the network planners because, without a valid overview of the network, it was difficult, if not impossible to plan for and route circuits through the network. In addition, it was difficult to forecast when a particular segment might start to have capacity carrying problems and to identify those segments that were underutilized.

Since the data from the Network Planning group was used throughout the firm, it was imperative that the data provided be correct and concise. Network usage data was used by management to make decision on whether to add additional equipment or to even begin work to light an additional fiber. Likewise, information from the Network Planning group was used by the Sales Team as both a point to sell additional services to existing customers and to point to areas, as determined by segment loading, where additional sales effort could be applied.

The implementation of an encompassing Network Planning Tool would have benefited the company in all aspects. The Network Planning Group would have been able to alert management to the issues that resulted in the problems detailed in Chapter Two prior to them occurring, would have been able to provide management with a viable set of alternatives without resorting to “trial-and-error” engineering, been able to provide specific and detailed utilization and capacity information to the Sales force tin order to allow for more targeted sales campaigns, and saved an untold number of hours to plan and route circuits through the network, through nodes, and to the final destinations. In addition, by being able to perform simulations (“what-if” games) with the network design, optimization strategies and possible critical bandwidth issues could have been identified and preventative measures taken without affecting the end user of the services.

In the end, the author was laid off as a result of the recommendation from his direct manager, as a result of the inability of the manager and the author to agree on the implementation of a Network Planning System. The direct manager was then released from his position (i. e. fired) in the following round of lay-offs and, in contrast to the author, was not the beneficiary of the severance plan. The author has, since this employment, been a consultant for the European Space Agency as an employee of a British firm and specializes in Telemetry and Telecommand Communication Systems, Configuration Management, and Spacecraft Data Simulators. The former colleagues of the author that are still working for the TSP in this case study are all either searching for new jobs, within or external to the firm. One is in India training the people who are to replace the Network Planning group. The direct manager of the author, to his best knowledge, is now working for a competitor to his former firm and one who has been rumored to be interested in buying the author's former firm.

References

Held, Gilbert. *High Speed Digital Transmission Networking – 2nd Edition.* John Wiley & Sons Ltd, 1999

Kartalopoulos, Stamatios. *Understanding SONET/SDH and ATM.* IEEE Press, 1999

Author's Note to Reference Section

The books mentioned above are the reference sources that were used by the Network Planning Team during this project and by the Author for the production of this paper. All content within this paper is based on the author's first-hand experience during the lifetime of this project from problem recognition to a test implementation of the final solution. Company-proprietary materials, including network diagrams and descriptions have been removed per request of the company and redrawn by hand by the Author so the Company can remain anonymous.