

High Reliability Switching:
*A Configurable Solution for Maximum
Flexibility*

A White Paper Produced by



Excel Switching Corporation
255 Independence Drive
Hyannis, MA 02601

Phone 508-862-3000
Fax 508-862-3155
<http://www.xl.com>

HIGH RELIABILITY SWITCHING: A CONFIGURABLE SOLUTION FOR MAXIMUM FLEXIBILITY

Abstract

Modular switching architectures that offer multiple, selectable high reliability options enable developers and service providers to meet the reliability requirements of their customers. Coupled with an open, scalable architecture, these systems deliver greater flexibility than ever before.

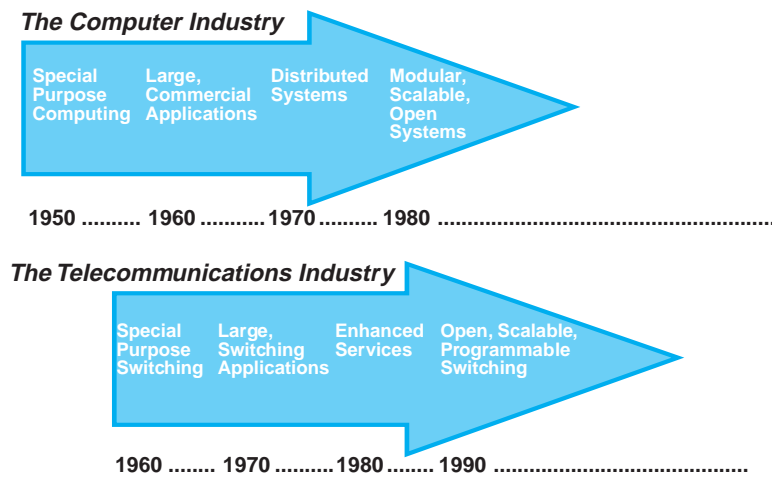
Contents

Section	Title	Page
<i>1.0</i>	<i>Introduction</i>	<i>1</i>
1.1	Today's Reliability Challenge	
1.2	Excel Switches: Designed for Flexibility	
<i>2.0</i>	<i>EXS: An Inherently Reliable Platform</i>	<i>4</i>
2.1	The Elements of High reliability	
2.2	Hot Removal and Insertion	
2.3	Chassis, Power, Cooling	
2.4	Internal Bus Architecture	
<i>3.0</i>	<i>Board-Level Redundancy Options</i>	<i>6</i>
3.1	Matrix CPU	
3.2	Service Resources	
3.2.1	MFDSP Card	
3.2.2	SRC Card	
3.3	Common Channel Signaling Interfaces	
3.3.1	ISDN PRI	
3.3.2	SS7	
3.4	Communications Line Cards	
3.5	Host Interface	
<i>4.0</i>	<i>EXS: Reliable Scalability</i>	<i>10</i>
4.1	EXNET Ring Reliability	
4.2	Nodal Isolation	
4.3	Controller and Ring Redundancy	
4.4	Service Resource Accessibility	
<i>5.0</i>	<i>PCX</i>	<i>13</i>
<i>6.0</i>	<i>Summary</i>	<i>13</i>
<i>Figure 6-1</i>	<i>Features, Advantages and Benefits</i>	<i>14</i>

1.0 Introduction

It's no secret that the telecommunications industry is undergoing dramatic growth in geographic reach, network infrastructure and service provisioning options. Concurrent with this growth is the inevitable evolution in technology that is needed to meet these changes. Today's service provider requires, indeed demands, increasingly greater flexibility in terms of the switching architectures that will address these new services.

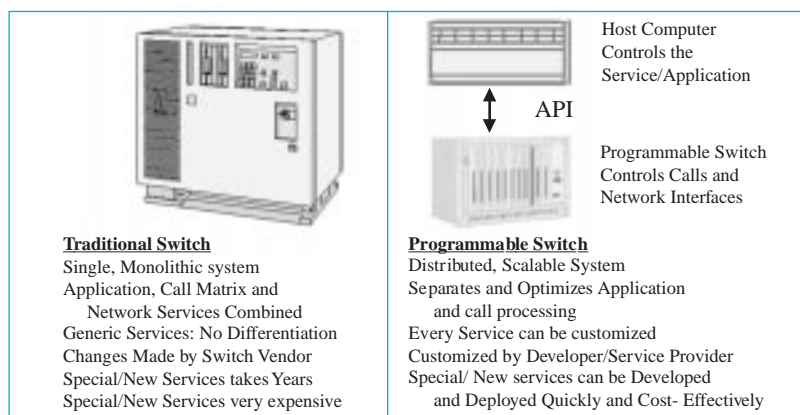
Figure 1-1
Comparison of Computer and Telecommunications Industries



Traditional switching systems that are monolithic in nature, and (by definition) centralized in practice, do not adapt easily to the dynamic changes taking place all around them. As a result, the switching industry faces the same technological evolution that the computer industry faced ten or fifteen years ago: the trend to move switching power and services further and further out into the network, creating scalable, distributed, client/server based telecommunications services. This trend will return the traditional switch to its original purpose of centralized regional switching, leaving distributed switching and enhanced services to the newer, more flexible architectures.

Early solutions to enable the quick delivery of enhanced services such as voice mail were developed on host computers that were connected to the local central office switch. As these enhanced services grew in popularity, the host could no longer support the switching requirements necessary to keep pace with the volume of calls. Enter the programmable switch, which operated as a front-end to the host computer, and a “back-end” to the traditional switch in the network. These switches offer the flexibility to turn up new services quickly, even when those applications require significantly increased call routing. Their compact size, efficient design and relatively low cost enabled them to be distributed throughout a network.

Figure 1-2
Comparison of Traditional Switches with Programmable Switches



Due to their inherent flexibility and scalability, these programmable switches are now evolving as the ideal solution to the challenge of distributing the switching infrastructure itself, and offering enhanced services on an as-needed basis in the network. Today's evolution of switching allows scalable, flexible, distributed programmable switches to be an integral part of an evolving integrated network and service infrastructure. Coupled with an open software architecture oriented toward flexibility and reliability, these systems feature modular code that is reused in many disparate applications and services.

1.1 Today's Reliability Challenge

Many standards have been established to help assure reliable communications in the network, such as NEBS in the United States, and CE Mark in the European Union. Meeting these standards is the first requirement toward network interoperability. Standards defining protocol-level compatibility for SS7 and ISDN signaling also specify reliability requirements. Yet other standards identify electrical and emissions compliance. These standards set the stage for interoperability and reliability in the network. Yet there is no standard to dictate switch reliability; switching manufacturers have the freedom to define the degree to which their products will address reliability.

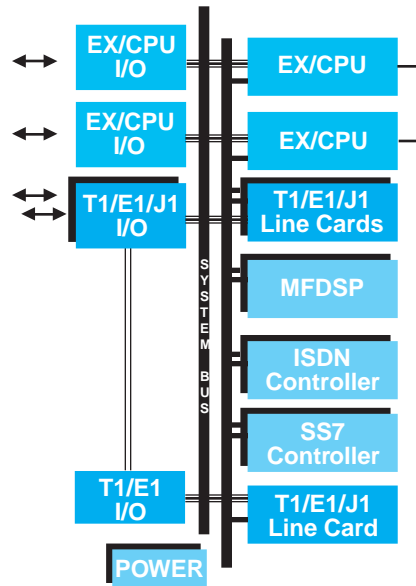
In order to be considered a reliable component of the network infrastructure, programmable switches must meet the same rigorous requirements that traditional switches have consistently delivered: they must be able to ensure continuous availability to their users. This translates into a level of reliability that will withstand system and component faults, and will continue providing service to the network even in the event of failure. Excel's solution can meet this challenge.

1.2 Excel Switches: Designed for Flexibility

Excel Programmable Switches were designed to address today's needs for telecommunications networks and services: an open, modular, scalable architecture that is fully programmable and for which the level of reliability can be selected to meet the service provider's needs. Designed with a dual internal bus architecture, all components of an Excel switch can be configured for reliability based on the service being offered. The figure on the following page illustrates

Excel's EXS internal design, including a dual bus architecture with redundancy options in all major component cards including the Matrix CPU, line cards, DSP resources, common channel signaling for ISDN and SS7, and other services.

Figure 1-3
A Redundant Switching Architecture: the Excel EXS



Of paramount importance when service providers are selecting a switching architecture are the reliability and throughput that is required of the system. If the system is scalable, then throughput can be guaranteed as the subscriber base grows. But there is a trade-off being made between cost and reliability based on the service being deployed.

An example of a mission-critical application is a service node application. When the switch is being deployed in the network providing call processing services, even a few moments of switch downtime or unavailability can add up to significant revenue loss as well as loss of customer confidence. The service provider cannot afford the system to be unavailable, whether it is the result of a failed component or because the system is undergoing maintenance or expansion. This application demands the highest levels of reliability.

For some other applications, there is a trade-off between cost and reliability that dictates the level of redundancy that the system needs. For example, a customer premise environment application such as electronic mail provides an important service very cost-effectively for the small business owner. In this case, the customer may need redundant power supplies in order to keep the system running in the event of a power supply failure. But to address his need for a cost-effective CPE solution, the customer may choose not to configure fully-redundant host computer and switching equipment.

The flexibility of Excel switches allows service providers to select the level of reliability necessary for the service being provided based on their knowledge of their business. Together with flexibility in terms of system size, scalability, and open programmability, the service provider has a platform that is configurable to meet his needs in every respect.

Excel offers four switching platforms. The LNX and CSN architectures use the same internal design, bus architecture, components and application programming interface. The LNX supports up to 2,048 ports in a twenty-slot chassis, and the CSN supports up to 1,024 ports in a seven-slot chassis. The EXS is a network architecture that enables multiple LNX and CSN “nodes” to be connected together delivering up to 30,000 non-blocking ports. The PCX provides a “solution in a box,” operating in a PC chassis and supporting up to 512 ports. Each of these products offers multiple levels of selectable reliability.

2.0 ***EXS: Inherently Reliable Platform***

The EXS platform address the needs of a Central Office switching environment. Using a common internal architecture, they are hardened to provide continuous 24-hour availability. This is achieved through a combination of hardware fault tolerance, software recovery and data integrity. This approach prevents failures from disabling the system, improves on-line repair, and accommodates online expansion without burdening the programmer or adding to system overhead.

2.1 **The Elements of High Reliability**

Excel has an established design goal to provide the highest level of reliability possible, and to achieve this without impacting the operation of the system and with minimal impact on the host programmer. Using a combination of redundant, load sharing, and “N+1” hardware techniques, Excel’s EX/CPU manages all components in the system. High reliability is achieved through three interrelated functions: fault detection, fault isolation and fault recovery.

Fault Detection

Self-checking logic is employed by the processors and system software to detect failures. In Excel’s distributed processing architecture, self-checking takes place at the component level. The EX/CPU also manages failure detection on the bus, as well as on the power supplies and cooling fans. Further, since transient errors can arise in the exchange of data between components, the EX/CPU uses protocol hardware timers to detect irregularities during data transfer.

Fault Isolation

Since error detection takes place at the component level, the EX/CPU can automatically disable the component in error. This protects the rest of the system from the failure, preserving the integrity of the other components and enabling the system to continue operating.

Fault Recovery

After a fault has been isolated, the system automatically invokes the appropriate recovery procedure, depending upon the component in error. For components that are configured for redundancy, the appropriate partner component will be placed in active mode. The system also generates an alarm that is passed to the host, and when appropriate, turns on an error indicator on the board that failed. This recovery process is discussed for each component in the sections to follow.

2.2 **Hot Removal and Insertion**

So far, the system has detected failures, isolated them from the rest of the system, and conducted the appropriate recovery procedure so that continuous operation can be maintained. The next step

is to replace the failed component so that the system can be returned to its original status. With hot removal and insertion, components can be replaced without interrupting system operation. Without this capability, the entire switch would have to be taken out of service: the system would be shut down, the replacement or expansion component installed, the system rebooted and the software reconfigured. This would apply to either replacement or expansion, and for the duration of that downtime, there would be no service, no revenue, and probably not many happy customers!

The Excel internal bus is designed to support hot removal and insertion of component cards. When a card is electrically detached from or inserted into the bus, a signal is generated and the system software can execute the appropriate action. When a component is inserted, the system software will conduct diagnostics, identify the board and initiate the appropriate action to bring it into service. This way, a component that has failed can be replaced while the system is in operation, without interrupting any other components in the system.

This “hot pluggability” is of paramount importance, not just for error detection and correction, but also during maintenance and expansion operations. If the service provider needs to add communications interfaces such as T1/E1, service resources such as digital signal processors (DSP), or common channel signaling resources such as SS7 or ISDN, these resources can be added to the system while it is running, without interruption in service.

2.3 Chassis, Power and Cooling

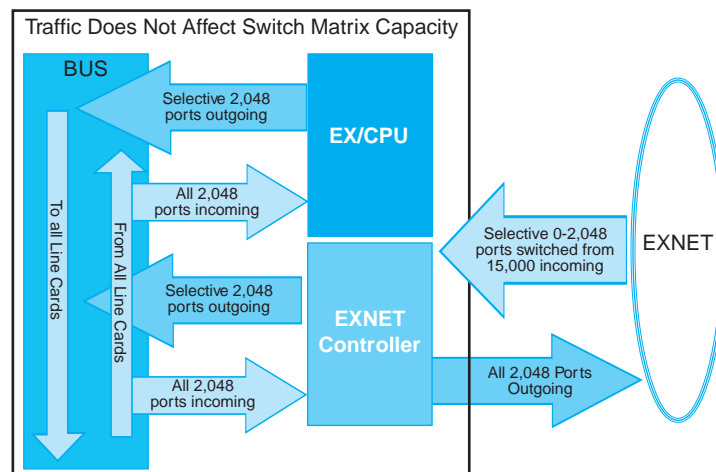
The EXS integrates, in a single system, multiple components including the matrix CPU, line and service resources in a chassis that was designed for reliability. The compact chassis are rack-mountable and are designed to operate in a standard office environment, without special air conditioning or other environmental controls. Emissions compliance and cooling are achieved without the need for a large cabinet enclosing the system components.

The chassis are designed with dual power input, and support optional dual power supplies. The system is actually designed so that one power supply can provide the power required to run a fully configured system. When dual components are used, the system load-shares the power; in the event of a failure, the remaining power supply will provide adequate power to the system. If a failure occurs, a message is delivered to the host, and an error indicator is lit on the front panel of the failed component. Similarly, there are two fan trays, each with multiple cooling fans and dual input from both power supplies. Each fan is individually sensed to monitor rotational speed, so that a single fan failure will create an alarm that is detected by the system software. In the event of a power supply failure, both fan trays will continue providing cooling to the chassis. Since both the power supplies and fan trays are hot-pluggable, a failed component can be replaced while the system is running without interrupting service.

2.4 Internal Bus Architecture

The EXS chassis uses a dual mid-plane bus architecture that supports the attachment of redundant EX/CPU cards and manages arbitration between them. The dual system bus provides a fast, high capacity connection among distributed processing elements. The system is designed using a selective space switching architecture. This enables all resources direct access to all ports in the system, maximizing the number of ports available for network connection and enhancing performance. The bus is designed to support efficient message transfer and linear growth. The figure below illustrates the selective space switching architecture.

Figure 2-1
Excel's Selective Space Switching Architecture



One bus supports the active EX/CPU while the second bus supports the backup EX/CPU. All other major cards are connected to both buses. In the event of an EX/CPU failure, the standby EX/CPU will take over processing, thus the second bus will take over bus operations.

3.0 Board-Level Redundancy Options

All board level components of the EXS are designed for high reliability, either by supporting redundancy or load sharing at the board level, or by allowing N+1 protection. Board level (1+1) redundancy is required for critical single-card reliability, such as the matrix CPU. N+1 reliability is a cost-effective approach when multiple cards of the same type are configured, and one card can serve as a backup in the event of a single-card failure. In a non-critical environment, any or all of these cards can be simplexed, giving the service provider the greatest flexibility in balancing system costs vs. the potential cost of failure.

3.1 Matrix CPU

The EX/CPU is the reliability “engine” for the entire system. It executes the system software that manages error detection, isolation and recovery for all components. In addition, each EX/CPU card is configured with two physical control interfaces to the host computer supporting a 10Mb Ethernet interface (see Host Interface, Section 3.5).

A system with redundant EX/CPU cards is configured with one EX/CPU operating as the active card. All call processing is managed by the active EX/CPU, while the standby card handles synchronization messages to remain in hot standby mode with the active card. In hot standby mode, the EX/CPU maintains a mirror copy of all connections and all system level configuration information. In the event that an error occurs on the active EX/CPU, processing will automatically be switched over to the backup so that the system remains in operation and retains all stable connections. An alarm is generated and passed to the host computer, an error indicator is set on the failed component, and all host-to-switch communications continues with the newly-activated EX/CPU.

With hot pluggability, the failed component can be replaced without interruption of service. When the failed card is replaced, the host computer can reconfigure the system to place the new card into active service. The alternate will then be placed back into standby mode, and the system will be restored to its redundant configuration.

In an application requiring high reliability, the system may also be configured with redundant host computers. In this case, each host will be connected to one of the EX/CPU. In the event of a host computer failure, the standby host can initiate switchover to the standby EX/CPU.

3.2 Service Resources

Excel service resources include the multi-function digital signal processing (MFDSP) cards supporting tone reception and generation, conferencing and other service resources and the substrate controller card (SRC) which enables the system to “bit” switch up to 8 calls per DS0.

3.2.1 DSP-ONE Card

The DSP-ONE card is designed to support multiple functions using the same hardware. Each card has its own timeslot interchange, and is configured with multiple digital signal processors and/or voice recorded announcement SIMMS that are then made available on a pooling basis to all calls. Multiple DSP cards can be configured in a single switch, enabling the service provider to configure as many DSP resources as are needed for his application. With multiple cards, specific DSP services (for example, conferencing) can be configured on two cards, so that if an error is detected on one card, the other resource can still handle that particular service. When an error is detected, the EX/CPU takes the failed card out of service, removes those DSP resources from the available pool, sends a message to the host, and sets the error indicator on the failed card. The switch automatically uses the alternate resources, which will continue handling the workload until a replacement card has been installed. With hot-pluggability, the failed card can be replaced without interrupting system operation. The host will be notified so that the replacement card can be configured to resume sharing the workload.

3.2.3 SRC Card

The SRC card is supported with card-level redundancy, allowing the system to switch the service over to the standby card in the event of a failure. Operating in hot standby mode, all connection information is mirrored on both SRC cards. If an error is detected, the EX/CPU takes the failed card out of service, and conducts automatic switchover without host intervention. All subrate connections will remain in service throughout the switchover process, ensuring continuous system availability. The system will also send a message to the host computer and set the error indicator on the failed card, so that the card can be replaced without interruption of service.

3.3 *Common Channel Signaling Interfaces*

3.3.1 ISDN PRI

ISDN redundancy is supported by redundant ISDN PRI cards connected by an ISDN Redundant I/O Card. The Redundant I/O card uses two slots in the rear of the midplane, connected to the two ISDN PRI cards in the respective adjacent slots on the front of the midplane. Using the selective space switching technology, both ISDN PRI cards have access to the same PRI D channels. In the event of a hardware failure of the active ISDN PRI card, the standby card automatically takes control of the signaling channels and continues call processing operations. Because there is an HDLC path connecting the two cards, all channels that are in a stable state are retained. The failure detection and recovery mechanisms are handled automatically by the system software, without host intervention.

3.3.2 SS7

Network-based services such as those offered in the Intelligent Network must be able to deliver continuous service. The failure or removal of an SS7 link cannot cause an interruption of service; it must be so resilient that even a single call will not be dropped. This is why Excel developed its own integrated SS7 product line, designed to capitalize on the robust architecture of the switch and the multi-level reliability of the system. Excel SS7 uses a combination of hardware redundancy, mirroring technology and sophisticated software to ensure that any single point of failure will not interrupt service. With built-in fault tolerance, reliability extends not only to all stable calls; it even protects calls during set-up and tear-down. Excel SS7 reliability is accomplished automatically, without host intervention or programming. SS7 reliability is achieved at the network level as well as in the switch itself.

Network Reliability

The SS7 standards dictate that any switch connecting into the SS7 network must maintain the integrity of signaling by ensuring that signals are not lost, replicated or out of sequence. A set of standard SS7 messages has been defined to assure signal integrity as well as routing and other network management functions; the challenge is to assure that the link is reliable. Excel achieves this by supporting redundant links within an associated link set. If an error is detected on a link, the SS7 card will reroute traffic over the redundant link. In addition, with quasi-associated signaling, Excel's SS7 can be configured with multiple link sets, to provide redundant signaling to multiple routes into the network.

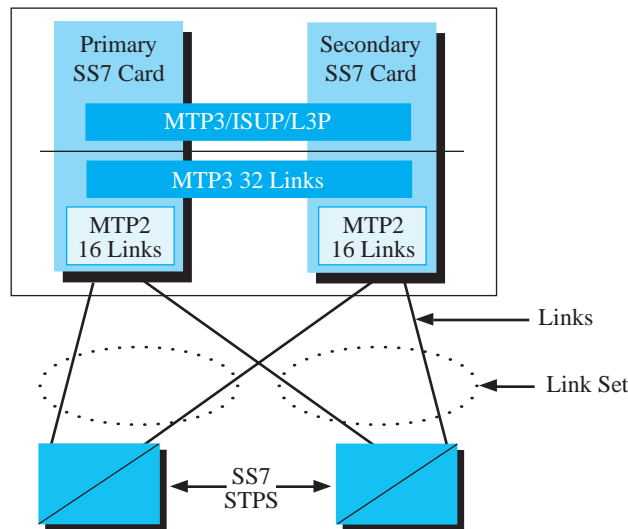
SS7 Hardware Reliability

SS7 redundancy consists of three components: two SS7 cards that are located in adjacent slots in the LNX or CSN chassis, and an SS7 Redundant I/O card, which occupies two slots in the rear of the midplane behind the respective SS7 cards. The SS7 Redundant I/O card connects directly to

both SS7 cards. Both SS7 cards have mirrored images of the MTP3, ISUP and L3P layers, thus both cards are aware of all signaling that is taking place. If one card fails, the EX/CPU will disable the failed card and will automatically switchover to the card that is still running. This implementation ensures that no signals are lost, not even those that are in process of set-up or tear-down. At the same time, the MTP2 layer that manages the link set is being executed simultaneously on each card, providing redundancy at the link level. Figure 3-1 illustrates the redundancy option for Excel's SS7 product, and how it appears to the external SS7 network.

The STP in the network sees only one Point Code, but communicates over links that are distributed to both SS7 cards. In addition, link sets can be split between two SS7 cards for reliability. This accomplishes two very important features: it provides link level redundancy and it allows the SS7 cards to do load sharing, which improves performance. The result is that the service provider is gaining performance out of both SS7 cards.

Figure 3-1
Excel's SS7 Redundancy Option



When switchover occurs, the EX/CPU notifies the host of the failure and sets the error indicator on the failed SS7 card. With hot-pluggability, a new card can be inserted and brought into service while the system is running, without interrupting service and without any host intervention or programming. Once inserted, the EX/CPU notifies the host, and at that time the host may reconfigure the MTP2 links to take advantage of the load-sharing feature of the SS7 design.

3.4 Communications Line Cards

Excel supports T1, E1, and J1 (the Japan variant of T1) digital communications. These cards can be configured to support an N+1 reliability option. For reliable communications the card set consists of a T1, E1, or J1 line card and a Redundant I/O card associated with each line card, plus one Standby I/O. When configuring for high reliability, the service provider will configure the number of card sets needed to support the subscriber population, plus one additional card set of the same type. If a card failure occurs, the EX/CPU will take the primary card out of service, and notify the host. The host can then reconfigure the line cards to route all calls connected to the

failed I/O card to the standby I/O card, which will reroute the call to the “plus 1” line card. This feature provides the ability to restore and switch to a fully operational state without physically servicing the switch.

3.5 Host Interface

Excel systems use an Ethernet I/O card for host communications, which is connected to the rear of the EX/CPU card. These two cards work together as a card set, providing high speed communications between the EX/CPU and the host application. They can be configured for redundancy by installing two card sets, each one including an EX/CPU and an Ethernet Controller. One card set acts as primary and active, while the other acts as backup. The backup card set keeps in synchronization with the primary via an HDLC link which continuously communicates status changes between the card sets. If either the EX/CPU or the Ethernet card fails, the system will automatically switch over to the alternate card set. The host will be notified, and an error indicator illuminated on the component in error. A replacement card can be inserted and brought into service without interrupting operations.

4.0 EXS: Reliable Scalability

EXS systems can grow to over 30,000 ports in capacity. Using Excel’s selective space switching design, EXS enables all ports to have direct access to and from all 30,000+ ports on the EXS ring, resulting in an expandable non-blocking architecture. And since these nodes can be added into the EXS network without interrupting service, the system is truly scalable. A service provider can begin with a single switch operation, and grow the network as needed, when needed, as subscriber populations grow.

It is critical that every element that makes up the expanded architecture be configured for redundancy, so that no single point of failure can cause the system to become unavailable. As a result, the links that connect the distributed switch shelves must also have a redundant architecture. To achieve this, the EXS provides three levels of reliability: Ring Reliability, Nodal Isolation and Controller/Ring Redundancy.

4.1 EXNET Ring Reliability

The EXNET Ring is actually a dual, counter-rotating fiber-optic network operating at 1.2 Gbps. Configured with one or more EXNET Controllers and I/O card sets in each node, the ring provides a fault-tolerant backbone to the system. Both rings connect into an EXNET I/O card in each node, and because each EXNET Controller has its own 2,048 by 15,360 port timeslot interchange, all nodes have access to all ports on the ring. One ring operates as the active ring, while the other operates in standby mode, handling synchronization and error-detection processes.

In a configuration with a single EXNET Controller (i.e., one dual ring), one segment of the ring is nominally selected as the “primary,” handling all traffic and completing a continuous circle through all nodes. The other segment of the ring sends traffic in a counter-rotating direction from the primary ring, and is handling synchronization and error detection processes. If an error is detected, the ring segment in error is isolated from the rest of the system and traffic is automatically rerouted (making a “U” turn) onto to the counter-rotating part of the ring. This process is called folding the ring. Traffic remains on the counter-rotating segment until it reaches

the segment that was isolated, “folding” back to the primary segment once again, and continuing around in a racetrack fashion. An alarm and indication will be set so that service can be initiated on the failed segment. Since the standby ring is continuously kept synchronized with the active one, processing continues without interruption and without the loss of any stable call, and no host application intervention is required. The concept of folding the ring is illustrated on the next page. The top illustration depicts the traffic on the dual rings prior to a failure; the bottom illustration depicts the traffic after a failure has been detected. The portion of the ring that had connected Nodes A and C has been isolated, and the traffic has been routed, in racetrack fashion, around the “folded” ring.

4.2 Nodal Isolation

With an EXNET ring, a new level of switch reliability can be achieved. With multiple nodes in an EXS ring, the system can achieve nodal high reliability. If one of the switches in an EXS ring becomes unavailable, (or if the EXNET Controller or I/O in a simplexed Controller configuration fails; see below) the ring will fold at the node in error, isolating the failed node from the rest of the ring, but continuing operation with the remaining switches. This is especially important when system maintenance is being conducted. The service provider can load new services or software releases, or can even add another node into the ring without disrupting the service that is being provided on the remaining nodes. Nodal isolation is accomplished without host application intervention.

4.3 Controller and Ring Redundancy

In addition to ring reliability, Controller reliability must be assured to complete the redundancy requirements. This is done by configuring redundant EXNET Controller and I/O card sets. In this case, the system would actually consist of two dual counter-rotating rings (yet another level of reliability). Each pair of rings connects into one of the EXNET I/O cards, which in turn, passes messages to its partner EXNET Controller card. Both pairs of rings are carrying all of the same traffic at the same time, while the Matrix CPU at the receive end determines which Controller will place timeslots on the midplane in each node. The redundant controller cards are handled the same way as other redundant components. With error detection, isolation and correction, the matrix CPU automatically reconfigures the standby EXNET Controller as the active one, without interruption of service and without host application intervention. All stable connections to the node are retained.

4.4 Service Resource Accessibility

In an EXS system, external service resources connected to one switching node can be accessed by any node on the ring. This provides another level of reliability to the switching environment. If one external resource (for example, a voice recognition device in a messaging system or a vocoder in a wireless network) becomes unavailable or overburdened, corresponding resources connected to other nodes in the ring can be used. The availability of external resources on the ring has other benefits: it provides a way to load-share voice and other resources, and it reduces system cost by eliminating the need for the resource devices to be connected to each switch.

Figure 4-1
Three Node Ring in Standard Configuration





Figure 4-2
Three Node Ring after Detection of Ring Segment Failure



5.0 *PCX*

Excel's PCX switches offer a unique solution for CTI applications, most often found in customer premise environments (CPE). The PCX "solution in a box" enables switching resources, host processing and voice resources to be co-located in a single PC Chassis. Since it is designed to deliver voice services to environments where the trade-off between cost and reliability is more sensitive, this platform provides reliability where it counts the most: in the power source and in software error detection and diagnostics.

Power Source

The system supports redundant power connection and power supplies that operate simultaneously during normal operations. If one of the components fails, the power source or supply that is unaffected has the capability to provide adequate power to the switching operation until a replacement part can be installed.

Error Detection and Diagnostics

The system software has sophisticated built-in error detection and software diagnostics, with appropriate messaging to the host. When a card failure is detected, the system will set the error indicator on the appropriate card.

6.0 *Summary*

Excel has made a commitment to meet the rigorous demands of today's emerging telecommunications industry by delivering switching platforms that offer the highest levels of flexibility found in the industry. Flexibility not only in terms of programmability, modularity, scalability and performance, but also in terms of reliability. The service provider can select the platform most suited for his application, and can select the level of reliability that meets the requirements of his constituency. And when true high reliability switching is needed, Excel meets the highest standards in the industry today.

Figure 6-1
Features, Advantages and Benefits of High Reliability Switching

Feature	Advantage	Benefit
Dual System Bus	Provides the underlying reliability architecture that supports all redundancy options in a single node or system	High reliability is provided as an intrinsic part of the system design, without sacrificing other design criteria such as scalability, high performance and open programmability
Dual Power/Cooling	Redundant power supplies and fan trays share the workload during normal operations, can handle the workload in the event of a failure, and can be replaced while the system is running.	Workload sharing puts less stress on these components, and in the event of a failure of either a power supply or a fan tray, the system will continue operating with no interruption of service.
Hot Insertion and Removal	All major components (boards) as well as power supplies and fan trays can be inserted or removed while the system is running, without impact on any system resource.	This supports not only failed component replacement, it also allows expansion of system resources without interrupting service, providing a “grow as business grows” expansion path.
Fault Detection, Isolation and Recovery	Internal hardware and software provide a method to detect failures and isolate them from the rest of the system, and recovery procedures ensure continuous system operation.	Hardware failure detection and isolation helps the operator identify and complete the appropriate maintenance activity, while recovery ensures that service is not interrupted.
Component Redundancy	<p>Critical components such as the matrix CPU can be redundantly configured; one acts as the primary while the other maintains synchronization in hot standby mode, and can take over active operation in the event of failure.</p> <p>Other critical services including ISDN and SS7 use similar approaches to provide redundancy at the component level, coupled with sophisticated software to ensure continuous operation with their respective networks.</p>	<p>The EX/CPU is the heart of the switch and reliability functions. The combination of redundancy and hot standby mode ensure that this critical system component will provide continuous service to the system’s end users.</p> <p>Both ISDN and SS7 must ensure network reliability, so that no call, whether in set-up or in progress, is lost (or duplicated). Excel’s ISDN and SS7 reliability ensure call integrity both in the switch and with their respective networks.</p>
N+1 Reliability	Applied to line cards, N+1 reliability allows one additional card (per type) to provide backup to all cards of the same type, so that if one card should fail, the backup will be able to take over its operation.	This provides a cost effective solution to systems supporting several T1 or E1 line cards. With one card as backup, the system will continue supporting the full complement of DSOs, providing continuous operation without performance loss.
Load-Sharing Redundancy	Services such as DSP services support a redundant load sharing environment in which the redundant card shares the workload during normal operations, yet will be able to handle the entire workload in the event of a component failure.	This solution provides maximum efficiency and use of the cards while still providing the protection that is needed for these services.
Ring Reliability	This is the ability for switching nodes to be connected together to provide scalability, while ensuring that the network, or ring, is just as reliable as the individual switching nodes.	This allows service providers to configure systems for today’s subscribership, and to grow or distribute systems based on customer demand, without interruption of the ongoing service. Provides the service provider with the same level of reliability and continuous operation in large systems as they get with the individual switches.



Excel Switching Corporation designs, manufactures, markets and supports a family of open architecture, programmable switches used in over 50 countries worldwide, in a wide variety of telecommunications services including wireless and wireline service provisioning, enhanced service platforms, Intelligent Networks, and Personal Communications Services.

*Excel Switching Corporation
Phone 508-862-3000*

*255 Independence Drive
Fax 508-862-3155*

*Hyannis, MA 02601
Web Page <http://www.xl.com>*

WP-0107-02