

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

ISP IMPROVEMENTS IMPLEMENTED IN TSP1, 2 AND 3

1	INTRODUCTION.....	2
1.1	SCOPE	2
1.2	ABBREVIATIONS.....	2
2	TSP 1	3
2.1	INTRA-NODE REDUNDANCY.....	3
2.2	INTER-NODE REDUNDANCY	4
2.3	FAULT MANAGEMENT.....	5
3	TSP 2	6
3.1	INTER-NODE REDUNDANCY UPDATES	6
3.2	RESOURCE LIMITS.....	8
3.3	FOUR IO PROCESSORS.....	9
3.4	SECURITY	9
3.5	O&M PROCEDURES	9
3.6	TELOBR MANAGER.....	11
3.7	USER DOCUMENTATION.....	11
4	TSP 3	11
4.1	IMPROVED INTER-NODE REDUNDANCY	11
4.2	TROUBLESHOOTING	12
4.3	DATABASE DISK LOG.....	12
4.4	APPLICATION LOG.....	12
4.5	FILE SYSTEM.....	13
4.6	SECURITY	13
4.7	OTHER DATABASE ENHANCEMENTS	14
4.8	UPGRADES	15
4.9	FAULT MANAGEMENT.....	16
4.10	OUTAGE HISTORY	16
4.11	CRASH COLLECTOR.....	16
4.12	TELOBR MANAGER.....	17
4.13	HW PLATFORM – GEM (NSP 3.0).....	17

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.	
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen	
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A
		File	

1 INTRODUCTION

1.1 SCOPE

The scope of this report is to describe which new and enhanced functions that have been implemented in the last TSP releases in order to improve In Service Performance (ISP) for systems out in the field.

In Service Performance (ISP) can be defined as 'service reliability according to customer requirements' and deals with questions regarding robustness, reliability and maintainability of Ericsson's delivered systems. ISP, or rather a lack of ISP, is measured in down time.

ISP is a very critical performance indicator and customers, and in some cases local governments, are starting to penalize operators for non-acceptable ISP.

Several different factors contribute to the In Service Performance of a system. In order to achieve good ISP, our system must be user-friendly, the operators must have sufficient O&M knowledge, the exchanges must be correctly dimensioned and a fundamentally high level of quality throughout the system is required. In addition to this, a fault tolerant, robust platform that can keep the system in operation under virtually all conditions is needed.

This report describes the ISP enhancements made in TSP 1, TSP 2 and TSP 3. A summary of the software and hardware versions for each TSP release is shown in the table below. The definition TSP1-3 used in this document is thus only a way to connect different releases of JAP, TelORB and HW.

TSP	JAP	TelORB	HW
TSP 1	1.1, 1.4	AD9.05	NSP 1.1 (Nora)
TSP 2	1.1H, 2.0	R13	NSP 2.1 (Carmencita)
TSP 3	3.0	R14	NSP 2.1 (Carmencita)
			NSP 3.0 (GEM)

1.2 ABBREVIATIONS

DBN	Database
DNS	Domain Name System
DUCK	Dicos UNIX Communication Kernel
EIN	Ericsson Infotech
FTP	File Transfer Protocol
GCC	GNU C Compiler
GRA	Geographical Redundancy Agent
GUI	Graphical User Interface
HLR	Home Location Register
IDL	Interface Definition Language
IDP	Internal Delivery Package
IRP	Integration Reference Point
ISP	In Service Performance

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No		File
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen		
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A	

JAP	Jambala Application Protocol
JPA	Jambala Protocol Adaptor
LM	Load Module
LND	Live Notification Display
MDP	Master Data Partition
MO	Managed Object
Net Red	Network Redundancy
NFS	Network File System
O&M	Operation and Maintenance
OA&M	Operation, Administration and Maintenance
OS	Operation System
PB	Primary Backup
PM	Primary Master
POT	Persistent Object Type
SB	Standby Backup
SCP	Service Control Point
SM	Standby Master
SNMPv3	Simple Network Management Protocol version 3
SS7	Signalling System Number 7
SSL v3	Secure Socket Layer version 3
STP	Signalling Transfer Point
U-Qtil	Quick Utility
3GPP	3rd Generation Partnership Project

2 TSP 1

This was the first release of TSP and no improvements were therefore implemented. However TSP itself includes a number of features which functions are to provide a system with good ISP for example by supporting both intra-node redundancy and inter-node redundancy.

2.1 INTRA-NODE REDUNDANCY

TSP provides replication mechanisms to all data stored on a node, that is information is replicated to at least one more processor in the cluster. This mechanism enables TSP to handle software and hardware faults in a seamless way to the service. In case of a crash on one processor due to software or hardware fault, the replicate data is instantaneously accessible on the processor where the replicate resides while TSP tries to recover from the fault. Different levels of recovery exist with increasing escalation factor. If the recovery is successful, TSP automatically updates the recovered processor with the latest information from it's replicate and is back to the original configuration. In case of a permanent fault on the processor, TSP removes the processor from the configuration and appoints new replicates to data from the faulty processor to other processors. This is called self-healing.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

Upgrade of both application and TSP software can be done on the node with no disturbance to provided service. The application can use the smooth software upgrade method that allows ongoing processes to transfer their internal states from the old to the new software version. For upgrades of driver type of software or even the operating system itself, where old and new versions of the software can not co-exist on the same machine, the redundancy upgrade method can be used. Here the software in one processor at a time is upgraded by first removing the application software from it and re-distribute it on the remaining processors in the cluster. The emptied processor is now upgraded with the new version of the driver or the OS and then the traffic is restored to its original distribution. The upgrade process automatically proceeds with the remaining processors to be updated. Since the re-configuration of the software can be handled without affecting the traffic this method will also be free of disturbance of the application.

TSP allows adding of processors in real-time without any disturbance to ongoing traffic. The distribution algorithms in TSP support this. When a new processor is added TSP re-calculates the distribution for the new configuration by spreading the load evenly on accessible processors. The same mechanism is used when removing a processor from the cluster.

OA&M provisioning through OMLite and JPA load share between both OA&M processors. If one OA&M processor goes down, the other is used. This is controlled by the OA&M Scheduler, which decides which OA&M processor in a node is Master or Backup.

2.2 INTER-NODE REDUNDANCY

Inter-Node Redundancy or Geographical Network Redundancy is the concept of hardware and software duplication in separate geographical locations. In terms of ISP, there is no single point of failure on the node level. This enables the standby TSP node to transparently perform the same functionality as the primary TSP node at any time. The Primary TSP node processes traffic while the other node is in standby mode. The Standby TSP node starts processing traffic only when the Primary node experiences a failure.

The Geographical Redundancy Agent (GRA) is responsible for monitoring the state of the node on different levels to determine if the node is available for SS7 traffic. The GRA runs on the SS7 stacks with an instance of the GRA on both the Hot-Standby and the Primary nodes. The Hot-Standby monitors its own node and marks it as available or not available. The instance running on the Primary node is responsible for triggering traffic re-routing when the Primary node is marked as unavailable. A traffic re-route is achieved by deactivating all link sets of the Primary node. The Signaling Transfer Point (STP) then redirects the traffic to the Hot-Standby.

The Scheduler decides which node is Primary and which is Standby, and which OA&M processor in each node is Master and Backup. There are two schedulers on each zone, one master scheduler on one IO processor and the standby scheduler on the other IO processor.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

The Scheduler ensures that:

- When one node is not responding, either because it is down or there is a network failure, the other node becomes automatically the Primary.
- When one OA&M processor is not responding in a node, the other OA&M processor becomes the Master one. The states of the OA&M processors in the other node do not change.
- When one node changes from Primary to Standby or vice-versa, the states of the OA&M processors of that specific node also change. They go respectively from PM/PB to SM/SB. However, even if the OA&M processors are on the Standby zone, they can still update the database in the Primary zone.
- There is a primary node at all times
- The primary node is always the node best suited to handle the current traffic

The schedulers within the same node maintain a heart beat between each other through a dual redundant network and with the Schedulers on the other node through an inter-node network connection. The heartbeat mechanism allows each Scheduler to have instantaneous knowledge of the state of each one of its peers and the status of the links of its own node. This information is used by the Scheduler to determine which state it should be in at all times.

2.3 FAULT MANAGEMENT

Fault Management improves ISP since it will be easier to detect and localize faults with the help of notifications and alarms. Also, it provides access to more faults and more relevant data regarding these faults.

A component in Fault Management, the Event Handler, maps notifications/alarms from the platform resources (hardware and software) to managed objects. The Notification/Alarm Collector can then forward this notification/alarm to the manager application in the customer's Operation System (OS) using event channels based on the CORBA Notification Service. It is also possible for the managing system to choose which events it wishes to view and those it wishes to screen out.

Event Management facility provides Integration Reference Point (IRP) interfaces. IRP (pre-3GPP) is an Ericsson-wide standard for ensuring interoperability between product specific applications and the types of generic applications.

Two Specialized GUIs are implemented to display current alarms and notifications on the system: The Live Alarm Display and The Live Notification Display. The Live Alarm Display displays basic information (timestamp, severity, class and object reporting alarm) and detailed information about an alarm (history of states, list of comments, name/value pairs), and updates

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

alarm information automatically. The Live Notification Display displays basic information (context, meaning, MO Class and instance of the object that reported alarm, timestamp) and detailed information about a notification (history of states, list of comments, name/value pairs) and updates notification information automatically.

3 TSP 2

3.1 INTER-NODE REDUNDANCY UPDATES

3.1.1 Partitioned Network Redundancy

TelORB R13 adds support for running Geographical Network Redundancy in partitioned load shared mode. New APIs are added to the TelORB API in order to allow applications to use Network Redundancy Partitions.

Partitioned Geographical network redundancy is where any given object/subscriber has one node as its primary and the other as its standby and the whole set of objects/subscribers are "partitioned" into two subsets, of which one subset, or partition has Zone 1 as its primary and Zone 2 as its standby, and the other partition has Zone 2 as its primary and Zone 1 as its standby. This ensures that both zones are capable to run traffic.

With this function there is however, a significant impact on backup time, transfer time and external-sync time. The times increase by between 50% and 100% depending on the size of the database and number of files per Mdp (Master data partition). The reason for this is of course that the number of Mdps will effectively have doubled.

3.1.2 Other Inter-Node Redundancy Updates

Other Geographical Redundancy updates that impact ISP are:

- Quick synchronization
- Transfer Time Decrease
- Smooth TakeOver
- New Notifications and Alarms
- Stand-Alone Mode of Redundancy

3.1.2.1 Quick synchronization

Quick synchronization is a new feature that is used within a remote synchronization. Quick synch is the method used to compare two databases and construct one large compensating transaction. This transaction will be applied on the standby node to compensate for any lost update, create, or delete operations. The logic of the remote synch is to first try a quick synch and if that fails an external synch is attempted.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

Quick synch is faster than the existing external synch, but it has a limit of how large differences between the databases it can manage. The external synch doesn't have this limit as it copies an entire backup of the primary database to the standby node. Quick synch decreases the time window for a double node failure, thus improving ISP.

3.1.2.2 Transfer Time Decrease

The backup transfer time has decreased due to the new method to transfer the backup to the standby node.

Previously, the process for transferring a backup was:

- backup is started on the primary zone
- once the backup is finished the backup is transferred to standby zone
- an external sync is performed
- the transaction backlog is cleared and NetRed backup is created

In TSP 2 the process was changed to

- backup is started on the primary zone
- once the backup has started the backup starts to get transferred to standby zone
- once the backup is finished, one last backup transfer is sent for integrity reasons and because most of the backup has already been transferred, this transfer is quicker
- an external sync is performed
- the transaction backlog is cleared and NetRed backup is created

The improvement is that the backup begins transferring while the backup is in progress. The backup and backup transfer is done in parallel instead of sequential.

Transfer Time Decrease decreases the time window for a double node failure, improving ISP. During manual reconciliation, the standby node's backup is cleared and a backup dump is transferred. Manual Reconciliation is the process involved in the re-creation of the database of the standby node using a copy of the database from the primary node. During this time, the standby node is unavailable and if the master node would go down during this period, traffic would not be processed. With Transfer Time Decrease feature, the backup dump is transferred quicker, hence reducing the time window that the standby node is unavailable and the vulnerability of double node failure.

3.1.2.3 Smooth TakeOver

Smooth TakeOver is introduced in TSP 2. A smooth takeover process avoids the need of a manual reconciliation of a zone that has just gone down and it therefore reduces the possibility of a double node failure. Instead, if possible, the standby zone will be able to accept transaction updates from the Primary node keeping it up to date and then switch back once recovered.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

3.1.2.4 New Notifications and Alarms

New alarms and notifications have been added for the scheduler in TSP 2. The new alarms and notifications deal with node, IO and remote synchronization status.

The new Node Status Monitor Heartbeat feature is also added. It uses a notification reporting mechanism to send out Heartbeat messages reporting Node Status. It is configurable to send heartbeats (notifications) from every minute to every 60 minutes or not at all.

The Scheduler reads the state information from the State Supervisor. The overall health of the node is then calculated from the state information. A notification object is created and then raised holding the status of the IO processors, TelORB, the stacks, and the node itself. This process repeats itself at each heart beat interval. The notifications are sent to the LND (Live Notification Display).

In terms of ISP, these new notifications, alarms and heartbeat monitor provide information on the overall health of the node(s).

3.1.2.5 Stand-Alone Mode of Redundancy

Stand-Alone Mode is a software maintenance procedure used to separate zones, divert traffic to a zone, or upgrade a zone.

Stand-Alone Mode is when the update channel is deactivated via software; this will cause no database transactions to go across the Inter-Zone Link. When the Update Channel is deactivated, the zone will go to stand-alone mode. When the Update Channel is activated, the zones will go to the correct status

As far as ISP is concerned, this feature is used for maintenance of the node and provides for faster upgrade/repair times.

3.2 RESOURCE LIMITS

3.2.1 Operating System Resources

In order to further increase the system's robustness to software faults, there is now a limit to the number of system resources available to an individual process instance.

System resources are resources in the operating system, such as timers, ports (on which dialogues are built), transactions, sockets etc.

The purpose of the limit is to detect and abort processes that erroneously allocate too many resources before the total number of resources in the system is exhausted. The default limit is currently 400 and is intended to be way above the normal needs so that most processes should not need to care at all.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

3.2.2 Database Resource Limits

3.2.2.1 Memory Usage

It is now possible to control how much of the memory on a particular processor that can be used by the database (in %). If this limit is reached, transactions that attempt to create objects will be aborted. There is also a new alarm that will be sent when usage of the database memory gets to a certain level.

The share of memory set aside for the database, the amount of this memory currently used, and the alarm level, are all settable/viewable in TelORBManager.

3.2.3 MasterRecordTable

The maximum number of objects per processor is set in TelORB INITIAL DATA using the DbdRecordTabMasterSize. New resource controls are introduced for this table, as exceeding the size of this table will result in system crashes. There is also an alarm that will be sent when the usage of the MasterRecordTable goes above a certain level.

These settings are NOT visible in the TelORBManager.

3.3 FOUR IO PROCESSORS

In TSP 2, two additional IO processors were added to the basic TSP hardware configuration bringing the total IO processors to four. This improves ISP for the increased capacity reduces the risk that the processors will be overloaded and also allows for better distribution of resources running IO tasks. Typically, TelORB processing will take place on IO1 and IO2, and application OAM activities (i.e. provisioning) on IO3 and IO4.

3.4 SECURITY

In TSP 2 password protection is added on the two daemons that can be used to diagnose faults in TelORB. These two daemons are the U-Qtil daemon used for getting status of processes, ports etc. and the TelORB-shell used to get various internal TelORB states, and look at database objects. This will ensure that no unauthorised person will get access to the node, hence improving ISP. It should be noted that the daemons should only be used for troubleshooting by experienced maintenance personnel and they should not be used by the end-customers.

3.5 O&M PROCEDURES

3.5.1 Optimization of Backup Functionality

The backup functionality has been improved so that the creation of bootstrap tables takes shorter time. This means that it takes less time to take a backup

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved		Kontr - Checked		Datum - Date	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22		Rev A	
				File	

which will make it possible to run the backup with shorter intervals to ensure that the backup data will be more up to date.

The backup mechanism optimisations in TSP 2 showed the following changes in characteristics. (Note - without Partitioned Network Redundancy)

TSP Version	Number of HLR Subscribers	Backup in Primary Zone (min)	Transfer (min)	External Sync (min)	Backup in standby zone (min)
TSP1	1 Million	19.28	14.56	8.35	29.03
TSP1	1.9 Million	24.45	24.41	11.35	48.19
TSP2	1 Million	10.59	16.15	6.34	10.05
TSP2	1.9 Million	13.48	24.20	9.36	14.06

As can be seen, there is a marked improvement in the backup time in both the primary and standby zone.

3.5.2 Improvements to EPCT for Upgrading

Improvements have been made to epct in order to disallow immediately upgrade packages that attempt to upgrade TelORB (and which therefore require a Redundancy Upgrade) and add processes and POTs to applications (and which therefore require Smooth Upgrade) at the same time.

Also epct and the TelORB runtime have been modified so that upgrade packages now not only include information on what version of Load Modules to upgrade TO, but also what version to upgrade FROM. If there is any mismatch between the version required by the upgrade package and the version in the system, the upgrade will fail.

In terms of ISP, this will solve a number of trouble reports and improve the handling of the system.

3.5.3 Multiple Processor Enable/Disable

It is often necessary or desirable to disable or enable more than one processor in a system. For example, during the upgrading of a "Nora" generation cabinet to a "Carmencita" generation cabinet, it will be required to disable and remove all old processor cards, and add and enable all new processor cards (not all at the same time).

When expanding systems due to increased subscriber numbers, multiple processors are also often added at once.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other) UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		Nr - No. UAB/R/L-01:082 Uen	
Dokansv/Godk - Doc respons/Approved UAB/UR/LC (Soli Liu-Viking)	Kontr - Checked	Datum - Date 2002-04-22	Rev A
		File	

Enabling or disabling processors can be a lengthy process in a large system since the Distribution Units in the system are re-spread out to get optimal load distribution.

In TSP 2, the capability to enable or disable multiple processors is added as a new O&M function. This will improve the handling of the system and thereby reduce the risk for faults being made.

3.6 TELORB MANAGER

From the first release of TelORB, there has been no supported O&M tool since it was agreed that our customers would provide their own. Due to the number of requests for such a tool, and due to requests for a supported high-quality tool, the TelORB Manager has been developed.

The handling of the system in service will be easier with the use of the TelORB Manager and ISP should thereby be improved.

Note that it is not obligatory to use the TelORB Manager. TelORB retains an open and published CORBA IDL interface for Operation & Maintenance.

3.7 USER DOCUMENTATION

In TSP 2, the User Documentation has gone through a major revision, which should improve the handling of the systems in service. The changes can be separated as follows:

- New structure for User Documentation based around either a list of all possible actions, or based around the TelORB Manager
- New detailed description of actions, based on the TelORB Manager
- Changes to existing procedures to accommodate changes to TelORB (e.g. add/remove processor)
- New procedure to describe things that have not previously been documented (e.g. installing patches)

Functional changes or additions to the documentation are marked with a special tag to make it easy to update your competence.

4 TSP 3

4.1 IMPROVED INTER-NODE REDUNDANCY

Some improvements have been made to the network redundancy functionality in TSP 3. The functionality previously included in the Jambala (JAP) "scheduler" has now been re-implemented and integrated into TelORB. The "scheduler" was a Jambala component residing on the IOs keeping track of the two zones and monitoring their states, as well as scheduling synchronizations and backups. The reason for moving the functionality is to

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.			
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

solve the problems related to lost contact between the IOs and TelORB and thus also the contact with the scheduler.

New functionality has also been added for the network redundancy interzone links. New alarms have been added to alert the operator to problems with the Links between zones in a Network Redundancy pair.

4.2 TROUBLESHOOTING

In order to ease troubleshooting of TelORB problems, "Error Codes" have been added to a number of Notifications, and detailed troubleshooting guides will be added to the TelORB User Documentation to allow troubleshooters to attempt to solve the problem themselves. This should ensure that the products that are released have fewer faults and also if a fault occurs in a live exchange it should be easier to solve.

4.3 DATABASE DISK LOG

The database within TSP (DBN), which already has extremely good availability, scalability and performance, now supports an even higher degree of persistency for selected data. It achieves this by logging changes to selected persistent data to hard disks. This can be used by applications to make sure that all updates to critical information, such as subscriber account balances in a prepaid system, will be recovered even if the entire zone should reload. This will improve ISP since critical information such as billing data always will be up to date.

Similar to network redundancy, persistent object types (POTs) can be configured (in runtime) to have disk persistence. A transaction that updates such data will store the changes in a log file on the support processors when the transaction commits. If the zone should reload then the active backup is loaded (as always) and after that the database log files will be read to restore all disk persistent data.

POTs can not be both netshared and have disk persistence. This limitation may be removed in future releases. Transactions should only open one kind of POTs (netshared, disk persistent or normal POTs), different kinds should not be mixed in the same transaction.

4.4 APPLICATION LOG

The new application log feature allows an application to send log-messages to a file on the IO-machines. It offers a generic logging capability that allows applications to persistently store any information that they want, and it subsequently offers a customer querying capability over these logs. The application log should be preferred before printouts to the console. Some of the advantages are:

- Just one log-file for the entire application, e.g. not one per TelORB processor.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No		10 (10)
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen		
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A	

- Log-messages are time-stamped.
- Careful usage of the log-levels can be used to filter unwanted printouts.
- There is also an option to duplicate the log-messages to the consoles.

A log-message is sent to the IO-machines that write the message to a log file. Different applications can tag their log messages to end up in different log files. Every log-message has a log-level or priority. A level threshold can be set at various places to filter unnecessary messages.

In terms of ISP, this Logging functionality will keep track of activities on the node and help to troubleshoot when application faults occur.

4.5 FILE SYSTEM

An entire new distributed filesystem is included in TSP 3. The file system is used by the operating system itself to perform file access operations. A direct impact of the new filesystem is that the backup process is faster.

The DUCK-server and the FileServerInterpreter continue to be supported for the transition period.

A number of troubleshooting printouts are added to the system log in case of lost contact with NFS servers.

4.6 SECURITY

Some security features are introduced in the TSP 3 release. As a result, TSP 3 is more secure against possible hacker attacks, and it provides better integrity checking and auditing capabilities, hence improving ISP.

Support of Individual User Accounts

Individual user accounts for TelORB-related maintenance activities performed by system administrators on the IO processors are supported.

Minimization of System Services

UNIX services that are unnecessary for the normal and intended functioning of TSP systems are disabled. Popular administrative functions like telnet, rlogin, rsh and ftp have been disabled in favour of the more secure SSH suite of tools like ssh, scp and sftp.

Authentication Before Initiating a Remote Session

All remote user connections performed on a TSP node require authentication. OAM users are authenticated through the OAM toolbox. The suite of Berkeley r-commands are disabled on all Solaris processors in favour of the more secure SSH suite of tools.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No		File
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen		
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A	

Warning Messages

All users starting a remote session are presented with a configurable message. This message is intended to legally protect node owners, by declaring that the person accessing the system agrees to the terms and conditions described in the message.

Support of Remote Session Encryption

Information between a client machine and the TSP node can be intercepted. To limit the risk involved with exposing sensitive or confidential information, encrypted communications are supported for remote operations performed on TSP systems.

The SSH v2 suite of tools are used to replace more vulnerable UNIX system administration tools like telnet, rsh, and rcp.

The Secure Socket Layer (SSL v3) can be used in OA&M provisioning activities.

Highly confidential system data, such as passwords and keys, are encrypted.

Integrity Checking Capability

Through the Tripwire v2 software tool, system administrators can verify the state of the TSP node's Solaris file systems, by comparing it against a known-good configuration. The report generated can be used to assess the validity of changes performed on the node.

System Logging and Audit Trail Capabilities

Event logging is enabled on TSP's Solaris processors. The UNIX syslog facility accepts records sent to it by syslog-aware applications and middle-ware components residing on TSP. Scripts or system administrators can also use the "logger" command to send records to the syslog facility.

Solaris auditing is activated during system initialization. This kernel-level logging feature records system calls managed by the operating system. Audit trails can be used as a record of activity, and can be analysed as part of a forensic investigation.

The TSP system has the capability to support on-line storage of security audit data. Node owners and system administrators are expected to harvest and archive these files as part of their regular activities. File rotation mechanisms have been introduced in TSP to limit the possibility of file system overflows.

4.7 OTHER DATABASE ENHANCEMENTS

4.7.1 Database Implementation

An internal table known as the Lazy Copy Table has been removed and another table called the Record Table has been made able to expand in runtime. Some epct initial data parameters has been removed due to this, the most interesting is the DdbRecordTabMasterSize. This means that, in TSP 3,

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.	
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen	
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A
		File	

you no longer have to do dimensional calculations of how many database object that are expected to reside on each processor.

ISP will be improved since previously some outages have happened when the DbdRecordTabMasterSize have been set to low.

4.7.2 New C++ compiler

A new version (gcc v2.95.3) of the C++ compiler is distributed with TSP 3. It has a number of advantages over the old compiler (gcc v2.7) and the new compiler should be used in order to eliminate bad C++ code.

ISP improvements can be made since the new compiler finds more faults compared to the old one, which means that the faults can be found and fixed before the application is put into service.

4.8 UPGRADES

4.8.1 Site Independent Upgrade Packages & Patches

In previous releases, patches to TSP have been sent out as a number of updated LoadModules (LMs) with newer versions. This has worked reasonably well, but has placed the burden on generating upgrade packages with TSPs customer organisations. It would have been easier, and perhaps less error prone to send out ready made upgrade packages as well as the LMs. In the past this has not been possible because upgrade packages contain a number of instructions on exactly what to upgrade, and are therefore dependant on which IDPs are present in a particular system, which differs from system to system. In TSP 3 there are two major changes to how this process will work.

Firstly, updated IDPs will be sent out instead of updated LMs so that the granularity of the upgraded items is bigger. The updated IDPs will contain a number of updated LMs, and a number of unchanged LMs. The version number of the updated LMs will have changed, as will the version of the IDP. When building an upgrade package from this patched IDP, only the LMs, which have changed, will be updated in the system, so this is only a change in the handling, rather than the upgrade itself.

Secondly, the Upgrade functionality in TSP, and the toolset for generating upgrade packages have been changed so that the actions in an upgrade package are dependant on the relevant IDP being present. This means that ready built upgrade packages can be sent out to be applied to customer sites. For example, a particular upgrade package may contain an upgrade to the Java Virtual Machine LM, which is contained in the Java IDP. If the upgrade package would be applied to a system that contains the Java IDP, then the Java VM LM would be upgraded, whereas if the upgrade package were to be applied to a system without the Java IDP, the Java VM LM upgrade would be ignored.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No		10 (10)	
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen			
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File	
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A		

From TSP 3 onwards, upgrade packages are generated automatically in this way. This should improve the handling of the systems in service and thereby also improving ISP.

4.9 FAULT MANAGEMENT

An ISP improvement to Fault Management is the support of additional protocols SNMPv3 and 3GPP IRP, which provide more access to faults. Also, with alarm filtering, alarms can be displayed according to the user's immediate needs, providing better access to faults.

4.10 OUTAGE HISTORY

A TSP 3 system will keep track of its own history of outages both of the zone itself, and even outages of a Network Redundancy Partition.

There are 3 types of outages that are recorded.

- Any outage of availability of a particular zone (e.g. Zone Reload) is recorded.
- Any potentially longer loss of data (due to the database returning back to its state at the last backup after a zone reload).
- Any outage of availability of a Geographical Network Redundancy partition due to both primary and standby zones being unavailable simultaneously.

All incidents of these kinds of outages are recorded and can be viewed via the TelORB Manager. Outage reports are kept for a certain lifetime, default 12 months, which is configurable via TelORBManager.

With this information it should be possible to analyze the reason for the outage and thus also suggest needed changes in order to avoid a similar outage in the future and thereby improve ISP.

4.11 CRASH COLLECTOR

The Crash Collector is a new function, which will run on, the support processors and which will collect up maintenance and troubleshooting information should a severe "event" happen in TSP.

The Crash Collector is effectively a daemon, which will run independently on both support processors. Every 5 minutes, it will take up and scan the dumps directory and check for new dumps that may have occurred. Should it detect a dump of a particular severity, it will collect up a number of files, which are relevant for troubleshooting.

Crash Collector archives will be cleaned automatically after 30 days.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.	
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen	
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A
		File	

4.12

TELORB MANAGER

The TelORB Manager has been extensively updated in TSP 3.
Enhancements include :

- **Upgrade** - The Upgrade panel now lists the upgrade history of the system, as well as giving the possibility to switch on or off automatic backups, or change the path from which upgrade packages are taken. One can also attempt to abort upgrades that are in progress.
- **Backup** - Allows the configuration of scheduled backups (and remove synchronisations), both periodic (every x hours/minutes), and fixed (at a certain time every day).
- **Processors** - The management of the processor clocks has been moved to the Time card.
- **DNS/NFS** - Time has been removed from the DNS card, and NFS added. NFS allows the viewing of current NFS mounted drives, their status, and also allows adding of new mounts by submitting File system commands.
- **Time** - The management of individual processor time is added to the Time panel, which is now separated from DNS.
- **History** - This new panel lists out the outage history of the system.
- **NetRed** - This panel comes up automatically should one connect to a TelORB Network Redundancy system.
- **DiskLog** - This panel comes up automatically should one connect to a TelORB system that has been configured with the transaction disk log software.

These changes will improve the handling of the system and thereby also improve ISP.

4.13

HW PLATFORM – GEM (NSP 3.0)

4.13.1

Scalable and High Available SS7

In TSP 3 with GEM (NSP 3.0) a new scalable SS7 solution based on the Ericsson SS7 stack from EIN will be introduced. This will also improve the availability of the SS7 stack since TSP provides replication mechanisms to all data stored on a node, that is information is replicated to at least one more processor in the cluster. As described earlier this mechanism enables TSP to handle software and hardware faults in a seamless way to the service. In case of a crash on one processor due to software or hardware fault, the replicate data is instantaneously accessible on the processor where the replicate resides while TSP tries to recover from the fault.

Uppgjord (även faktaansvarig om annan) - Prepared (also subject responsible if other)		Nr - No.		
UAB/UR/L Kerstin Nordin, LMC/JJ/E Giuseppe Conte		UAB/R/L-01:082 Uen		
Dokansv/Godk - Doc respons/Approved	Kontr - Checked	Datum - Date	Rev	File
UAB/UR/LC (Soli Liu-Viking)		2002-04-22	A	

4.13.2

Power

HW faults are a minor part of the faults leading to ISP problems nevertheless some improvements have been made which might improve ISP for systems in service.

The power supply has been changed with GEM so that each card has its own power supply. Previously there was one power supply for all cards.