1.3.2 Recovery

Although elements of a Disaster Recovery procedure were in place, and the task of rebuilding the server and reloading tables has been considered, it has never been tested. It was already diagnosed that any such rebuild will face obstacles, which will take precious time to resolve, at that very moment when there is none. Additionally, no scripts existed to rebuild either the databases or the tables. Further, the activation of less onerous Disaster Recovery scenarios, such as straight database recovery, would also meet with problems that could be prevented. In summary, physical recovery could be done with some effort, but logical recovery could not be contemplated.

It was prudent to run a full rebuild in test circumstances, so that all such issues are provoked, identified and resolved. Disaster Recovery procedures can be improved as per the results of the exercise, and the executive can be provided with confidence that, in the event that a Disaster Recovery scenario was activated, the recovery would be uneventful.

1.3.3 Data Quality & Integrity

It was already diagnosed that the Workflow application, the content of the server (databases, tables, indices) was of very poor quality: the bulk of the data not been indexed; substantial amounts of the data was (unintentionally) duplicated; the same data was located in more than one place; etc. This is the cause of some of the application problems and data inconsistencies, which result in data contention and Production downtime. These issues would present the most time-consuming obstacles in the event of activation of a full Disaster Recovery scenario (1.3.2).

When the Production server was first built, 100gb of space on the HP/XP SAN was allocated to it. Although somewhat more than half was actually used, once data had been de-duplicated, keeping growth space in mind, the would be less than half. The HP/XP SAN is expensive, and this space could be released back to Infrastructure.

1.4 Purpose

It was prudent therefore to conduct a full rebuild of the Workflow Production server. The following major tasks were planned. Although speed would be gained, that was secondary, the primary goal was application stability and recoverability:

- correct all recovery-related issues (speed as well as integrity)
- rebuild all the databases physically and logically
- rebuild all tables and indices
- correct and improve indices where they existed
- introduce indices where there were none

Software Gems is providing *Sybase* support services for the Workflow servers, and was commissioned to execute the assignment on a Fixed Price, Fixed Quality (result) basis. In so doing, all peripheral issues that were identified were also resolved, such that the server and the databases therein, were upgraded to a minimum level of both recoverabilty and server performance.

1.5 Constraint

The components which formed the *OpenImage* application were not changed logically, as that application cannot be changed. However, the related databases were rebuilt physically.

The main constraint in the exercise was that the work done must not affect the Workflow application, as that would affect Production. Although the application could be changed, that was not a budgeted or planned task.

1.6 Exclusion

Although the existence and configuration of the offsite Disaster Recovery server was taken into account, documentation was not available and it was included in the scope of this work and therefore hard recovery issues are not documented here. Nevertheless, recovery procedures are now greatly simplified and recovery times are greatly reduced. All changed scripts need to be implemented, and the Disaster Recovery server needs to be be fully tested.

1.7 Result

The assignment was completed without anything significant to report. There were some delays due to Customer staff not being available as per plan. A small amount of downtime occurred during the first week, which could have been avoided by more robust testing on the Customer side (testing was within Customer's scope of work).

The Workflow Production server is now recoverable under a variety of threats, and the executive can be assured that full rebuilds can be executed without event, substantially faster and in a scripted manner [1]. Ordinary database recovery is also much faster and databases can now be transferred between the Production and Test servers without losing integrity or security.

The Data Quality & Integrity issues have also been been addressed, somewhat beyond the Quoted scope of work. Ninety nine percent of the data duplication has been removed, all indices that were identified (as limited by availability of Customer staff) have been implemented.

Any single access to the *Sybase* server is between ten and one hundred times faster than before. This is readily apparent in the reports. However, this will not be apparent to users of the the *OpenImage* application, as the parallelism that the *Sybase* server is capable of has been severely limited, in order to avoid deadlocks in that application (which is due to its primitive coding).

The data devices on the HP/XP San was reduced from 100gb to 48 gb. Within that, the available space for growth is 14% in the SAN, and 11 to 45% in the databases (each database has different growth expectations, 26% average).

The previous assignment, a partial upgrade the *Sybase* server to secure it, has now been completed to a respectable level. The server and the databases therein have now been upgraded to a level where support can be provided with confidence and predictability, and reactive support requirements are minimised.

---

1 The scripts produced in the exercise need to be integrated into the Disaster Recovery suite.

2 Change

This chapter provides an overview of the changes achieved by rebuilding the Production Server. As requested by the executive, the information is presented in a before-and-after style, for comparison purposes. The various relevant resources of the server, the use of which was planned to be reduced and improved, are compared. This is not an exhaustive list of resource comparisons, rather the selection of resources are compared, to provide a summary picture.

In order to allow for instant visual comparison, each pair (before and after) of graphs or diagrams on any page, is rendered in the same scale, and comparable components are presented in the same colour: thus an object can be readily compared (in terms of size and proportion) with the object of the same colour in the paired graph.

2.1 SAN Allocation

The Allocation of disk space on the HP/XP Storage Array Network, which is expensive, was reduced from 100gb to 48gb, as depicted in section 2.2. The reduction in the number of units additionally eases maintenance, object placement and monitoring issues.

2.2 Device Allocation

A *Sybase* Device was created for each SAN Logical Volume. The usage (Database allocation) within the *Sybase* Devices (SAN Logical Volumes) was improved as follows. There was little change to the Log and System disks, these are not shown:

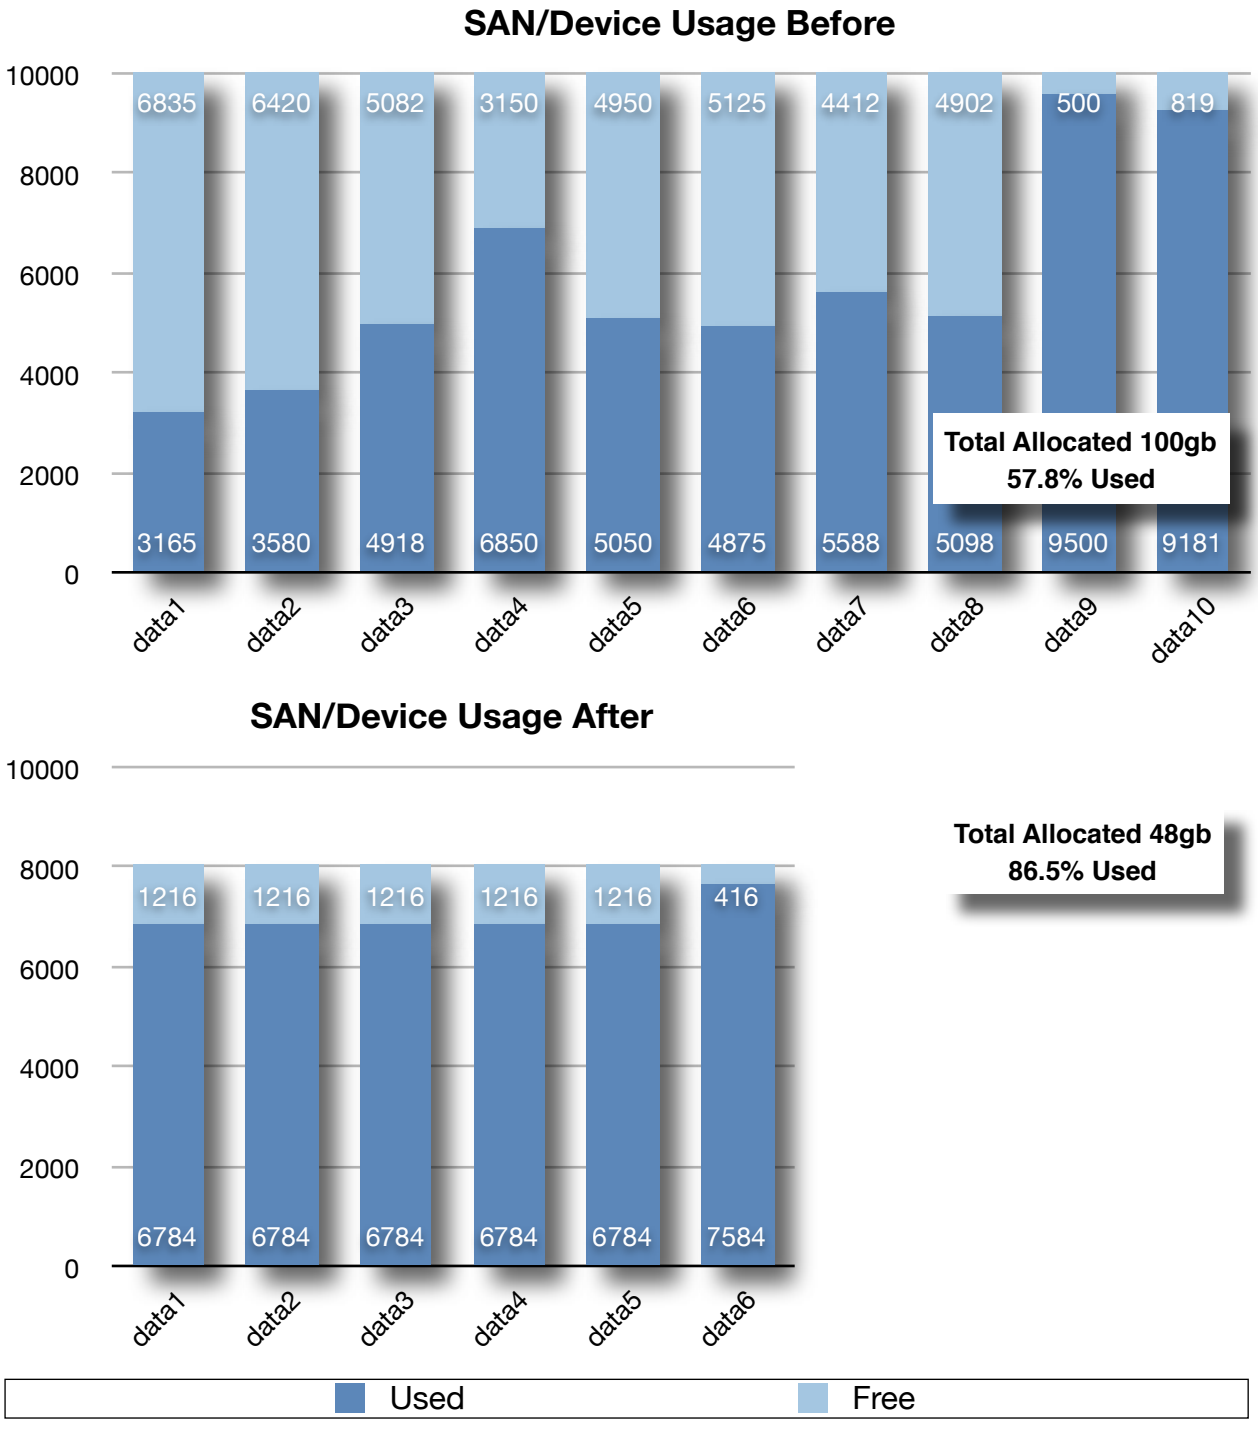

2.3 Databases

The number of databases was reduced, and the allocated sizes were corrected. This included amalgamations of databases, as well as merging of current and historical data from separate databases into a single unit. This eliminates the burden of people and application code having to search two databases for an item, and greatly reduces the overall number of objects and databases, which in turn improves speed. The following table identifies the database content, the allocated size before and after, and action taken.

Database	Content	Data/Log Before MB	Data/Log After MB	Action
`fai`	FAI (Customer) data storage, beyond the scope of *OpenImage*	11,924 242	12,288 200	
`polhistfai`	Copy of entire fai database containing archived Policy History	7,924 242		Eliminated The archived data was merged into `fai` Db (a large overlap was excluded)
`faidcdb`	*OpenImage* Setup for FAI (Customer)	15,231 506	5,376 300	
`corres`	Data storage for PAS (Oracle) images	13,711 153	23,040 100	
`archived_corres`	Archived data storage for PAS (Oracle) images	7,000 400		Eliminated The archived data was merged into `corres` Db (substantial duplicates were removed)
`FlexiNet`	Control values and staging area for IntraNet reports	135 12	200 20	
`bconnect`	Control values and staging area for TomCat	100 100	100 20	
`oai`	***Please fill in ...***	500 50	200 50	
`oidb`	*OpenImage* Reference	15 7	100 20	
`pas`	Control values and staging area for transfers to/from PAS	25 35	100 20	
`store`	Control values and staging area for transfers to/from other Customer systems	90 15	100 20	
`workdb`	Scratch area	1,000 400		Eliminated
`tempdb`	Internal Db for temporary tables	100 50	300 50	
`tempdb_MON`	Temporary tables for monitoring via MDA tables	90		Eliminated (Monitoring via *Sybase* Monitor Client is recommended instead)
`tempdb_intra`	Temporary tables for IntraNet reports	90	100	

2.4 Database Allocation

The space usage within the allocated Databases was improved as follows. This section also provides an overview of the actual data content that was reshuffled, ie. the result of the amalgamation, the archive merge, and the de-duplication operations. The small databases are summarised:

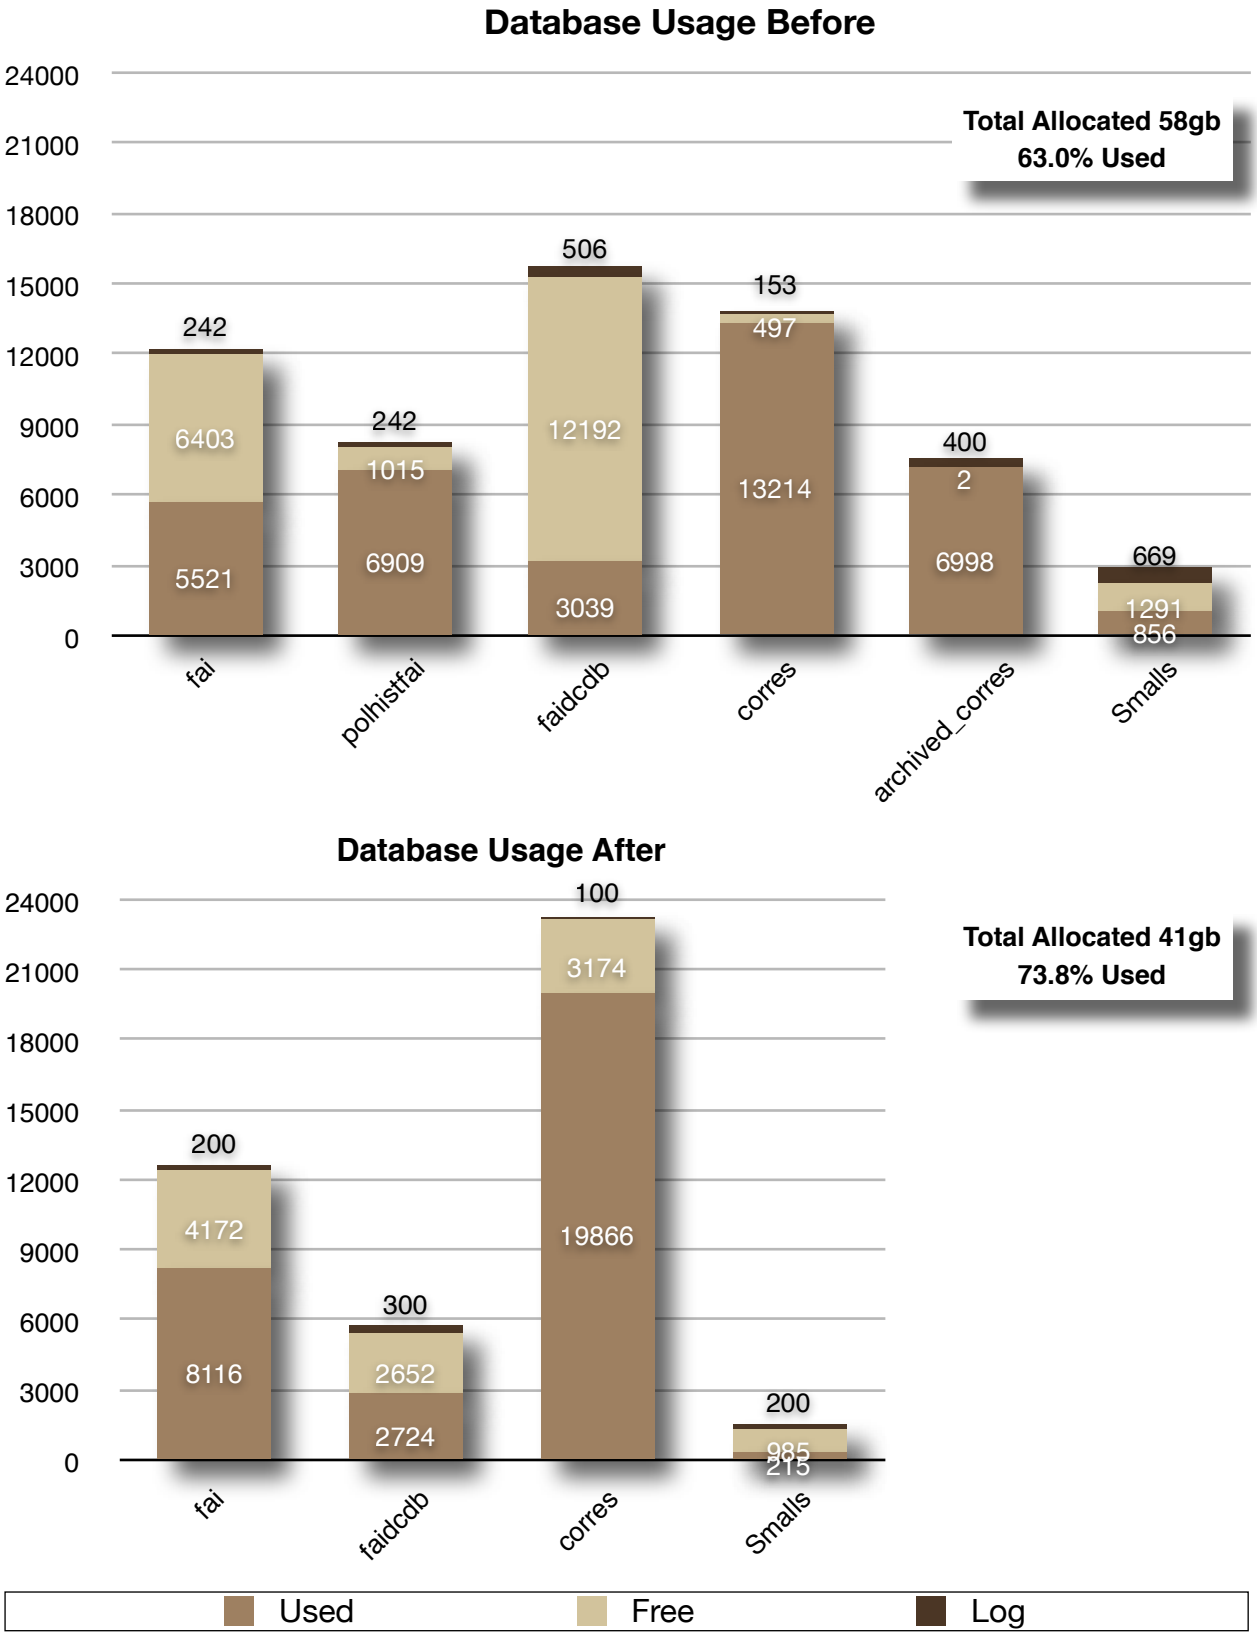

2.5 Index Improvement

The improvement (correction or addition) of indexing on the tables in each of the databases was the most important and the most time-consuming part of the assignment. It was known that the lack of proper indexing allowed substantial data duplication and adversely affected both speed and lock contention. As the Workflow application could not be changed or improved, the indexing changes within `fai` and `faidcdb` databases had to be carefully considered and tested. The database amalgamation allowed the overall number of tables to be reduced. This section provides a meaningful comparison of the reduction is overall database objects (tables and index types), as well as an overview of the corrected indexing.

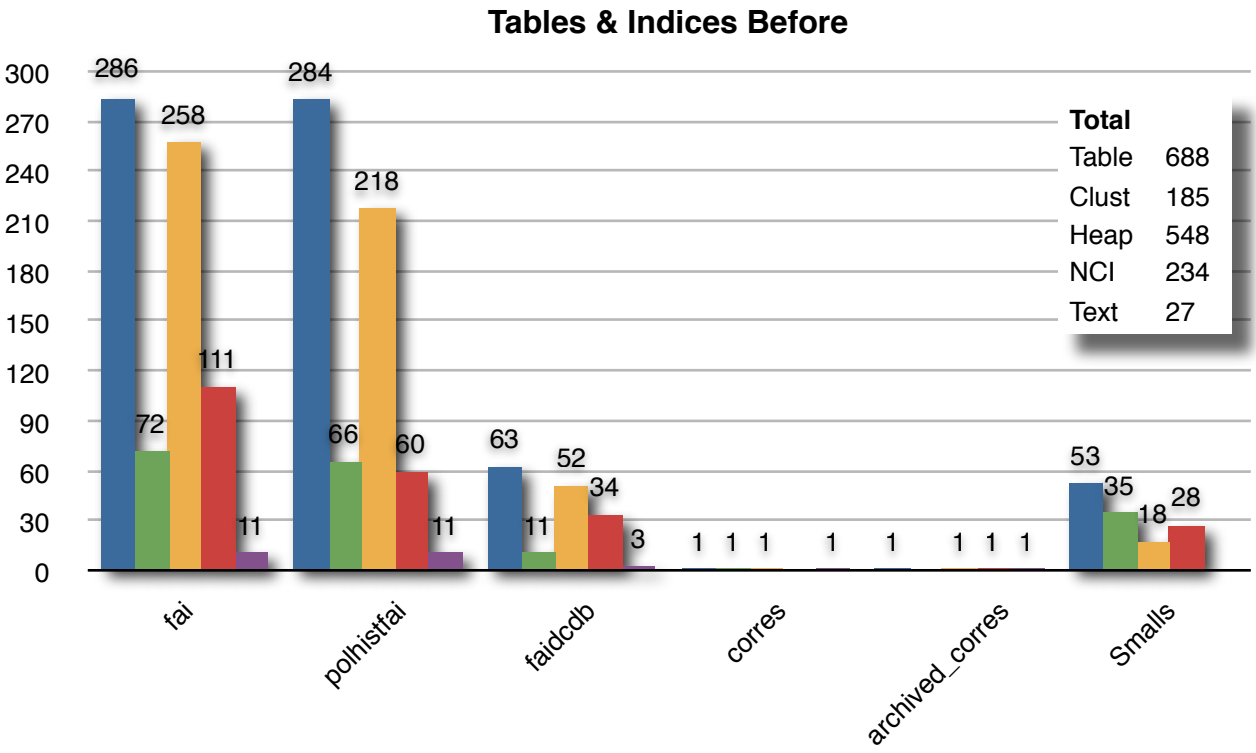

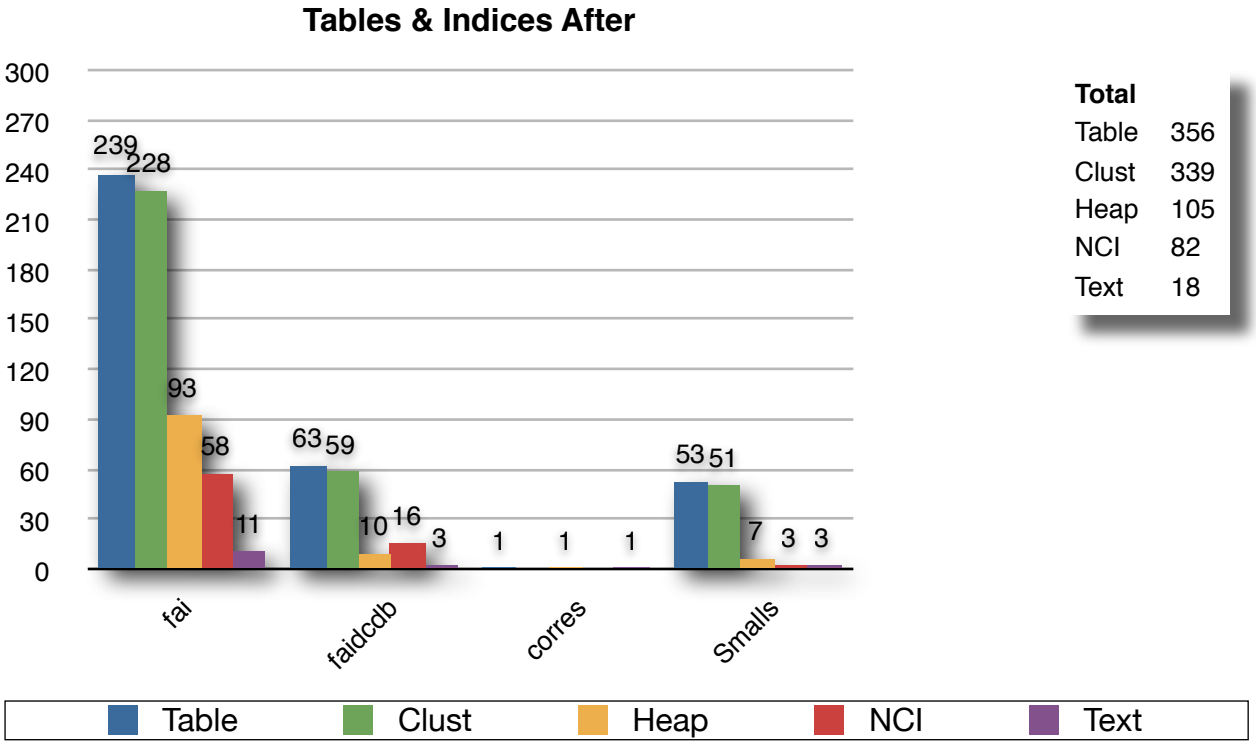

2.6 Index Addition

Every table in a relational database should have an index, and one of them should be unique. In a *Sybase* environment, one of these Indices should be Clustered. Unindexed tables cause serious maintenance and reload problems. Unindexed tables were corrected and Clustered indices were nominated where there was none (limited by information supplied by Customer staff).

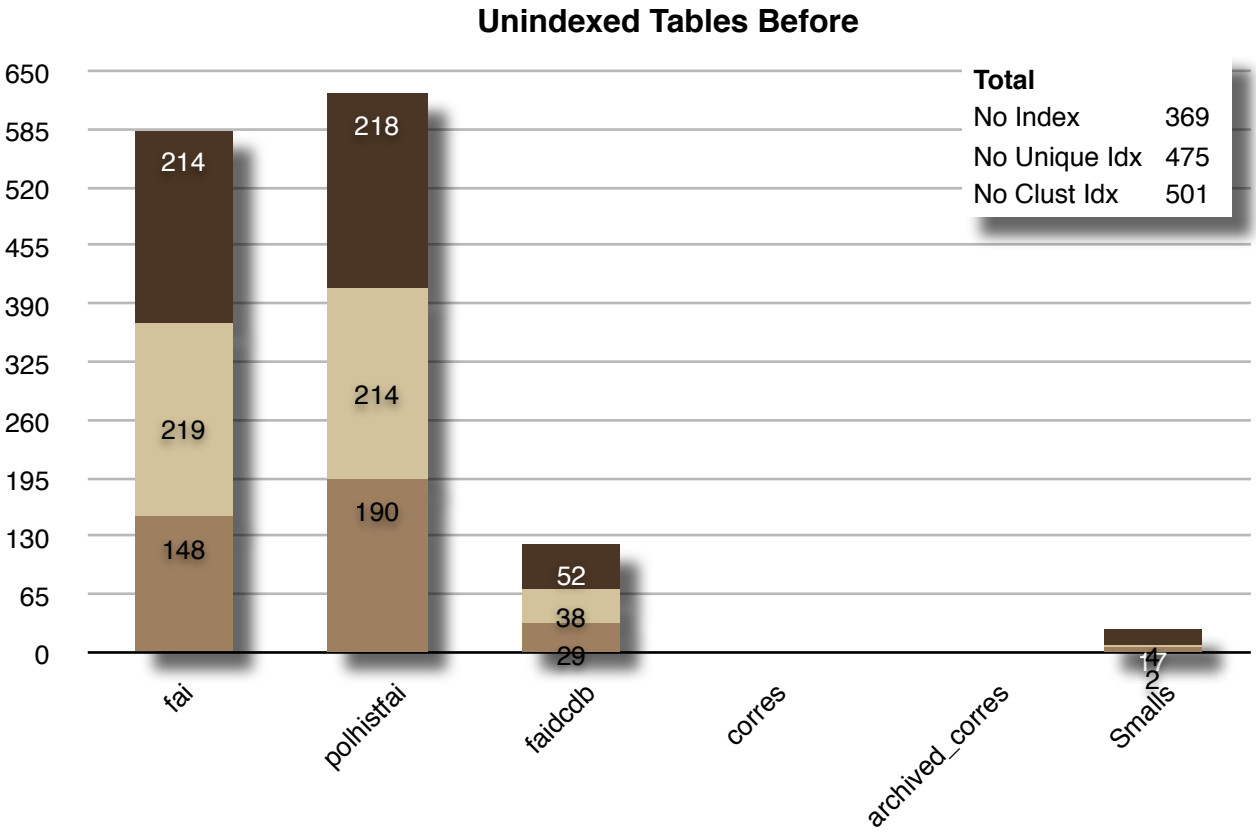

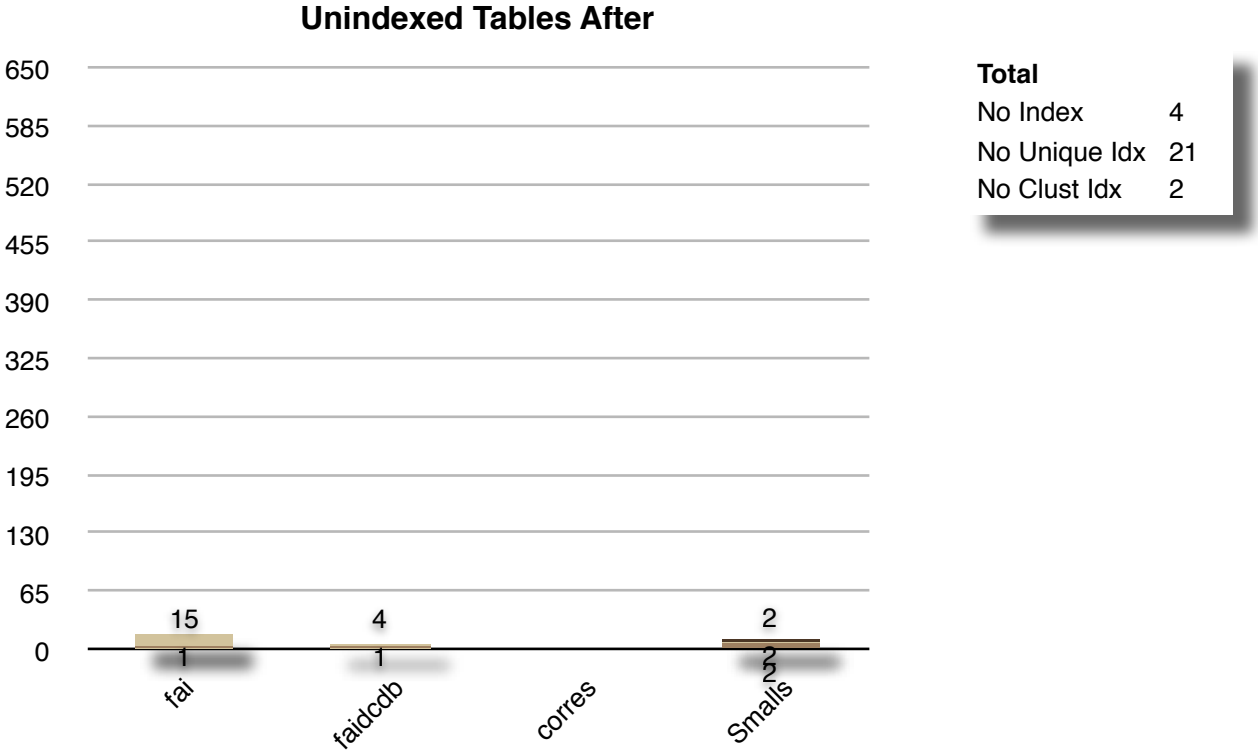

3.4 Rebuild Process

Refer to the diagram on page 3-5. Six complete iterations of the Rebuild process, including relevant testing by Software Gems and Customer staff, were executed on the Test Server. The main objectives of the iterations was to progressively improve the scripts, and to test again from scratch: this was essential due to the removal of substantial amounts of duplicate data, and the implementation of indices.

① Configure Test Server as per intended (final) Production Configuration
Test and progressively improve [2]

② Transfer Out Production
All databases, all tables, into shared `/nfs` disk area
Using Unload scripts (which invoke a series of *Sybase* `bcp out`)
Test and progressively improve

③ Transfer In Test
All databases, all tables, into shared `/nfs` disk area
Using Load scripts (which invoke a series of *Sybase* `bcp in`) with new physical defns
Test and progressively improve [3]

⑥ Build new indices (resolve and move duplicates) and Test [3]
Test Application against database [3]
(Allows for single tables to be unloaded, reloaded and retested)

The removal of duplicate data was also scripted, and these were executed in Production as per decision by Customer staff [4]. This allowed the next iteration of the rebuild to exclude the issues raised in the previous iteration, and to allow new issues to be raised in the next iteration.

The final rebuild on the Production Server was conducted in a single weekend.

② Transfer Out Production
All databases, all tables, into shared `/nfs` disk area
Using tested Unload scripts

④ Rebuild Production Server as per Step 1
Drop 10x10gb SAN Volumes; Create 6x8gb Volumes
Create *Sybase* Devices & Databases from scripts

⑤ Transfer In Production
All databases, all tables, from shared `/nfs` disk area
Using tested Load scripts, with new physical definitions

⑥ Build new indices using tested scripts, and Test [3]
Test Application against database [4]
Resolve issues

---

2 As planned, all Testing of the Rebuild process, and internal verification that the rebuild worked as intended, for each iteration, as well as after changes or corrections within each iteration, was performed by Software Gems.

3 As planned, testing of the application (ie. that the application worked unchanged) was performed by Customer staff, in concert with Software Gems.

4 Except for the Easter weekend, the duration of the Production server Rebuild, Customer staff were responsible for the Production server.

4.7 Server Overview

This identifies all relevant components of the Workflow Production server, and the DBA Desktop used to administer it and monitor its performance. For administrative purposes, since the Test server is administered from the same location, it is included.

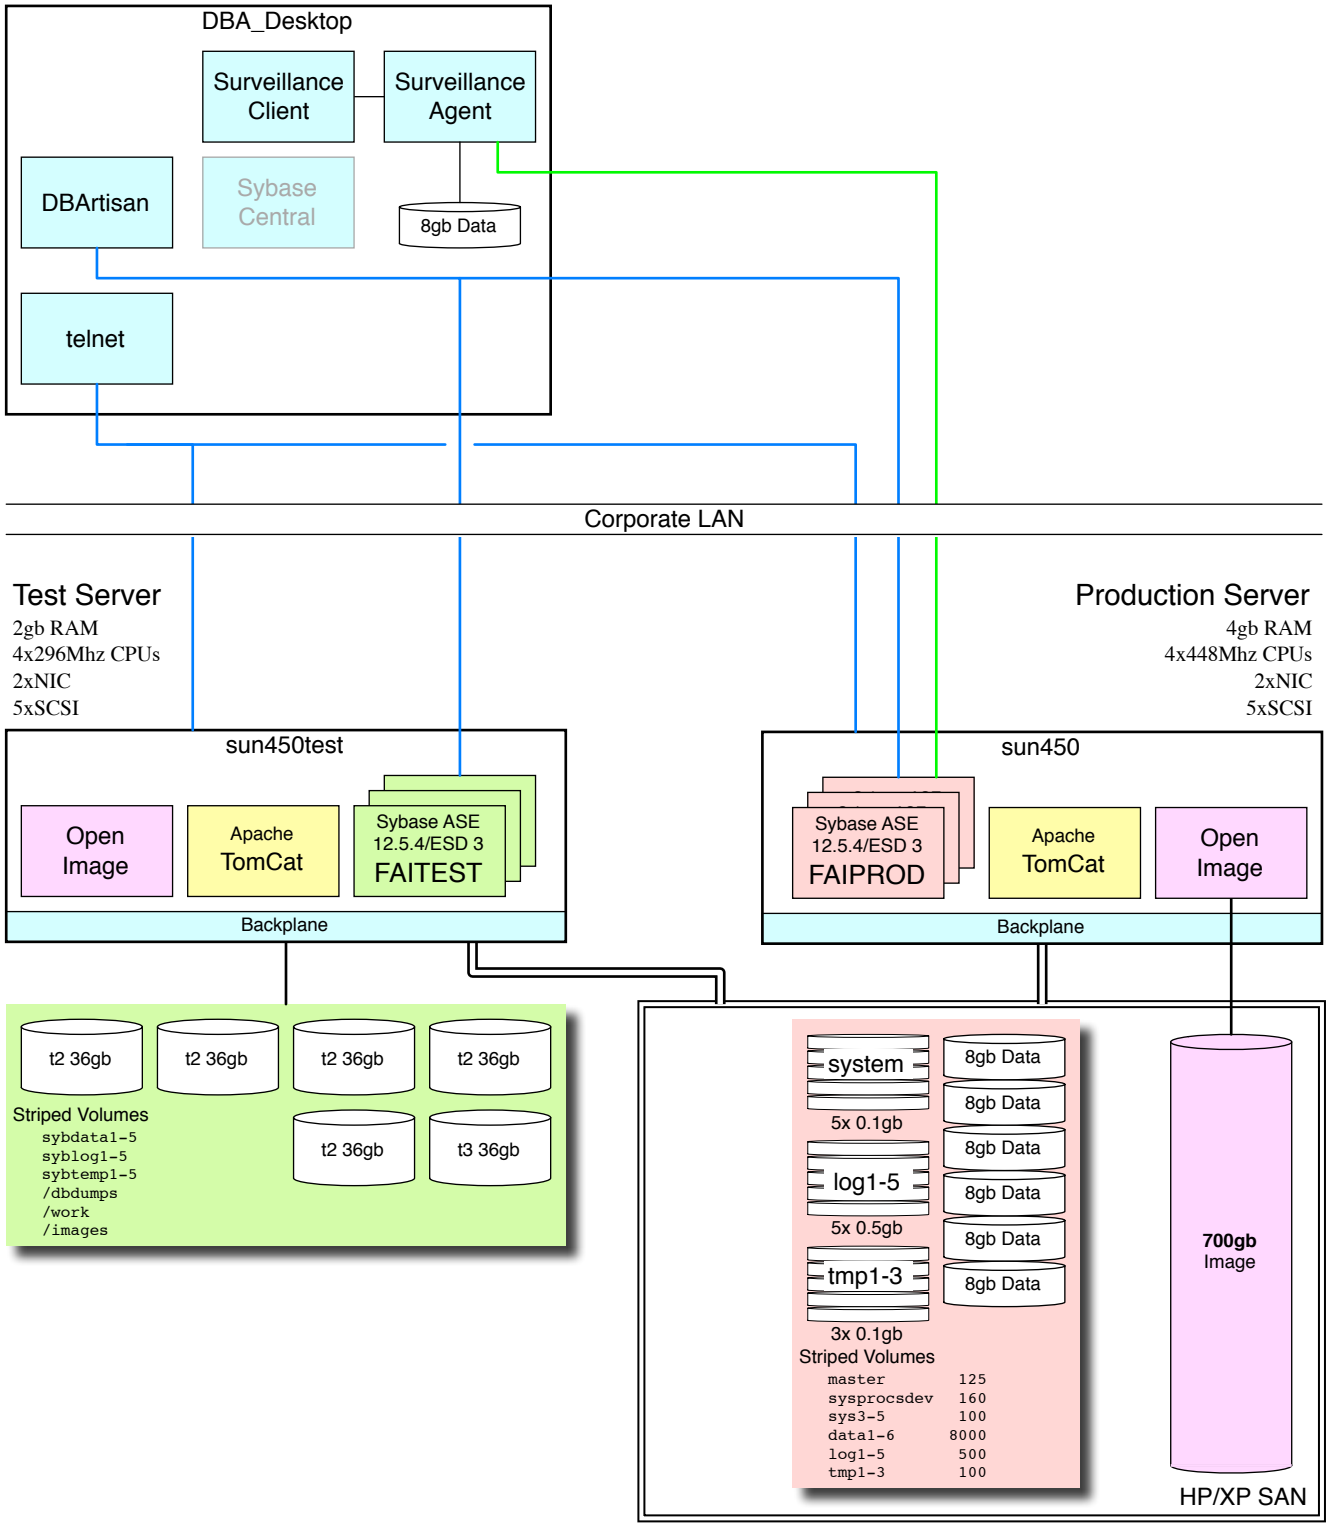

C.1 Application Overview