

Little Lessons From History

Chapter 8: New technology futility

One of the many pleasant pastimes I have indulged in during my TPF career has been to follow the fortunes of the many attempts to build replacement reservations systems to eliminate or emulate what TPF with PARS can do. There have undoubtedly been many more attempts than I know about, but equally undoubtedly none of them successful.

What these optimists are trying to do I call “NT RES”, because it always involves a so-called “new technology software platform” which is supposed to make life much easier and application development much faster. Looking back over 30 years of software development, it is difficult to conclude that these platforms are anything more than passing fashions. Anyway, as far as I can tell, there has only been one indisputably successful attempt at NT RES and that was Unisys’ USAS system. This was done in the early 1970’s, and was in many ways a copy of PARS into a different hardware environment (the UNIVAC mainframes of the day). USAS ended up with about 30% of the global market, primarily because of their strong foothold in Europe (Air France, Finnair, Iberia, Lufthansa and SAS), and IBM has kept the rest of the market for three decades now, thanks to TPF and PARS. Since the mid-1980’s nobody in his right mind would consider USAS as “new technology”, any more than TPF is, and application development in that environment suffers from exactly the same problems and restrictions that TPF has.

In 1975, when I arrived at KLM, there was something going on called NIBS (Neutral Industry Booking System). That was sponsored by IATA and supported by the travel agency community who had no direct access to any system at all at the time. NIBS kept a number of consultants supplied with Champagne and Caviar for some considerable time, but it never got

to a serious stage as far as I could tell/remember. I.E. it never got to the stage of a system software platform choice, let alone any application development. Hence, it is debatable whether one could classify this as a failed NT RES attempt, or as the germ of an idea which Sabre later exploited in offering travel agents direct connections to reservations functions. The latter in turn led to the GDS’s, as was covered in the previous chapter.

For several years in the late 1970’s there was a very significant effort going on called PRP (PARS Revision Project) and this involved nearly all the major TPF players of the time. IBM co-ordinated it and IATA was involved, but most of the work was done by American, BA, Eastern, KLM, Swissair, TWA and United. Even back in those days it was recognised that TPF’s structure was a problem constraining application development; the “system” was a rapidly expanding bowl of assembler spaghetti and things could only get worse. The idea was to redevelop the application from scratch using a high level language on a standard IBM mainframe operating system base.

In those days that meant PL/1 running under MVT with IMS or CICS as the transaction monitor (i.e. the new technology fashion of the time!). Many tens of man-years of effort went into the functional specifications and high-level design, but when it got to the stage of generating the first global estimates of the total development effort everything came unstuck. The project team came to the conclusion that the most appropriate unit to represent the development estimates was man-centuries. When senior management heard that, they decided they had more pressing things to do with their time and money, so they pulled everyone back to base. Nevertheless, one very good thing came out of PRP: IBM finally

recognised the crucial importance of a database manager capability for TPF and agreed to the adoption of the Swissair ACPDB package as standard, now known as TPFDF. We were thankful for small mercies.

The next NT RES attempt was in the early 1980's and was different. The manufacturer of non-stop hardware, Tandem (no longer with us, but in those days at the forefront of the non-stop platform fashion), entered into a joint venture with Lynjeflug (the domestic Scandinavian airline later taken over and submerged into SAS) and Bedford Associates (later taken over by BA/Speedwing) to build STARS. Although it nearly bankrupted both Tandem and Bedford, the system actually did get built and went operational in Lynjeflug. The owner trio touted the product around Europe, but it was functionally totally inadequate for anything other than a simple domestic airline environment, so nobody was interested. Then SAS took over Lynjeflug and that was the end of that. However, the rights to that software were sold to a group in Florida, called Phoenix-Air Inc. I believe, and 15 years later (would you believe) it turns up in China Eastern running on Stratus non-stop machines. Whether it ever went live in China I do not know, but I suspect not.

We have arrived in the mid-1980's, so now here is a very interesting one, actually two: IBM backing two horses at the same time. In Denver IBM Domestic was working with United to define an NT RES system based on VM (the obvious NT platform choice at the time?) and simultaneously IBM World Trade was working with KLM and SAS in Copenhagen to design an NT RES system under IMS Fast Path (the other obvious NT platform choice at the time?). TERESA, the one in Copenhagen, I knew all about since KLM (the "TPF will be dead in 5 years" clique), was involved. However, ENTERPRISE, the one in Denver, I only found out about more than a decade later when working with some EDS people who had been involved with it.

Both of those projects suffered the same fate as PRP. Once the development effort estimates were being generated and the man-centuries started piling up, economic reality takes over and everything gets canned.

Now we have arrived in the late 1980's and at least one thing is clear. The NT software platform choice is no longer an issue: it's UNIX. Proprietary mainframe architectures are dead and Client-Server with portable, open UNIX doing the grunt-work in the servers will rule the universe from here to eternity, right? However, there was someone who obviously had not got that message: Stratus. They came up with this brilliant idea to build from scratch a specially tailored system for Shanghai Airlines in China running under the Stratus proprietary non-stop operating system. The only restriction was that the end user interaction with the application had to look exactly like the system they had been using up till now: the complex command line entries of the CAAC USAS system. The system was built, I believe, and may have worked for the very special and simple Chinese domestic operation that carrier had, but on the other hand it may not have. It faded into oblivion like all its predecessors.

The real UNIX launch came from Israel with the AMSYS system. It was the product of a research project at the University of Tel Aviv involving Arkia Airlines, the "other Israeli carrier" and the one that is allowed to fly on a Saturday since it is not state owned. The structure of the system was the same as the existing TPF and USAS mainframes, except that the mainframe was a UNIX machine. For small, simple carriers who do not interact much with the airline world around them, it seemed to do an acceptable job for some time. It peaked at around 5-6 users in the mid 1990's, I believe, including such world-renowned airlines as Arkia, Croatian and LAP (the domestic Portuguese carrier now absorbed into TAP), but the base has shrunk since then. Also, in the mid 1990's the mighty and infallible IBM

purchased the rights to use a copy of AMSYS software to re-engineer it into an NT RES system. They did this as a reaction to what they saw as a serious threat: ReserVision (see below) was looking like it might be a success. However, after a couple of years and considerable amounts of money, plus the fact that ReserVision was obviously floundering, IBM decided to cut their losses and dumped it.

A whole host of similar developments based on UNIX sprung up continuously through the 1990's, but most of them were short-lived failures. In many cases they were being developed for start-up airlines who thought they were getting something on the cheap; usually the airline itself went bankrupt before the NT system got off the ground. The only ones I want to mention are those that still survive OpenSkies (which is NOT UNIX) and AirKiosk, plus the one I was involved with personally for several years: ReserVision.

Let's start with ReserVision, since it is a fascinating example of entrepreneurial skill. Although now dead and buried, in one respect it was a resounding success: it earned the small group of Swiss people who conceived it (unfortunately, not me), and spent all their time and energy and talents on it for a number of years, many millions of dollars.

ReserVision (RV) was conceived as a multi-tier, client/server, UNIX based reservations and departure control system using a relational database, which was to knock the socks off TPF with PARS. Development started early in 1992 by a company based in Zurich and early in 1994 the beta test version was taken to Bangkok and installed for Bangkok Airways. It became operational late that year and ran for about 18 months before the customer gave up and went multi-hosting in the Aer Lingus IPARS system. In early 1995 the product was installed for a start-up airline, Western Pacific Airways, in Colorado Springs. It ran there for nearly 2 years, then the customer decided to go multi-hosting in Sabre

(shortly thereafter they went bankrupt and disappeared). In mid-1996 the company owning RV was bought by Siemens AG in Germany (the Swiss went laughing all the way to the bank). Late in 1997, the product went operational at Hapag-Lloyd in Hanover, early in 1998 for Lufthansa Consulting in Cologne (for hosting a few small German domestic carriers) and later in 1998 for Bearskin Airlines in London, all very small scheduled operations. Even at that time RV was still struggling to implement fundamental things like ticketing, advance seat selection, GDS connections, change notifications, code sharing and many other capabilities, which all the classical systems had one to two decades previously. Since then, the classical systems have moved on to O&D revenue management, Internet access, customer databases, electronic ticketing and smart cards, none of which RV had even started to tackle. Application development was done by an Indian company in Mumbai and was a disorganised and undisciplined mess. The Indians proved conclusively that one can make a mass of unmaintainable spaghetti software much faster with VisualBasic than one can do so with Assembler. In addition, assuming a reasonably simple operating environment and constant tuning, RV could at best provide the performance throughput necessary to handle a 1-2 million passenger per annum airline. Performance in terms of system availability (i.e. keeping the system on the air) was another story. An operational RV system had up-times that made a bad TPF test system look like a model of stability.

By the spring of 1999 all 3 ReserVision systems in operation had been abandoned. Hapag-Lloyd and Lufthansa Consulting had moved to hosting in the Lufthansa mainframe system (USAS) and Bearskin moved into Speedwing's RTB mainframe (TPF). Siemens stopped supporting and promoting the product and decided to sell it. They never found a buyer as far as I am aware.

The only 2 (sort of) new technology systems which

are still alive and kicking to my knowledge are OpenSkies and AirKiosk. OpenSkies is not UNIX but runs on proprietary HP software, whereas AirKiosk is true blue UNIX. However, since they are still alive, I think it wiser not to comment too much on them. Suffice it to say that I do not believe that either of them will displace any of the big TPF or USAS systems, though it is possible that a few of the systems at the lower end of the transactions per second scale could be displaced.

In addition, there are several new efforts under way: Amadeus' NEW GEN, Iberia/UNISYS's AIRCORE & Sabre/Compaq's PINNACLE. If any of those look promising, then it can only be PINNACLE. However, we shall probably have to wait till after the end of my working lifetime before we can proclaim them successes or failures. Note, of course, that management never advertise their failures, so all you are likely to hear is total silence.

The bottom line of all the above: many hundreds of millions of dollars of airline money down the drain over a period of 25 years and not a thing to show for it, except in the balance sheets of some well known consulting companies and the private accounts of some lucky individuals. Despite this sombre fact, airline management, egged on by such consultants (i.e. those charging much more than \$200 per hour plus expenses), persist with the illusion that a new technology reservations system is possible if only those bigoted, anachronistic and antagonistic TPF people would allow it to happen.

Actually, it is technically possible, at least as a replacement for the smaller TPF systems, but management are going to drop dead when they see the price tag for building it and die a second time when they see the operational cost of running it. In return they will (possibly) get quicker application development, though nowhere near as quick as they think, and a system availability that only drops from

99.9x% of total time down to 9y.x%, where y, with a bit of luck, might be 9, but may be less and x will definitely be much less than 9.

Hindsight gives 20-20 vision, but no hindsight means no vision, which is another way of saying "he who does not learn from history is doomed to repeat it". Complex problems have no simple, easy to understand answers; they only have simple, easy to understand WRONG answers. For 25 years all the "TPF Killers", particularly those announced with a fanfare of publicity, have been abject and expensive failures.

*Bruce Taylor – Amsterdam
February 2002*