

Big, Bang, Boom Revisited: Why Large Projects Fail, A case study research of NPAC

Prof. Jim JOHNSON,
and
Prof. Dr. Hans MULDER,
University of Antwerp Management School,
Antwerp, Belgium

It was 7:30 AM on Tuesday, July 17, 2018, Sally Brown, was having a cup of coffee and watching “Morning Joe.” Suddenly she felt severe pain in her shoulders, neck, and jaw and a deep tightness in her chest. She picked up the phone to call Massachusetts General Hospital but the call couldn’t get through.

Later after Sally gasped for her last breath, “Morning Joe’s” Mika Brzezinski announced some breaking news: a key element of the US telephone system had crashed. Thousands of consumers could not receive calls. Brzezinski went on to say that the FCC and FBI were investigating the cause, but that it might be several hours or even days before the telephone system would be back to normal. Speculation on the show went into a frenzy, Was it the Russians? The Chinese? A disaffected hacker with a thirst for vengeance?

Later Sally’s family would learn the reason she could not get life-saving help during her heart attack was a cascading database error in the Number Portability Administration Center (NPAC) system and database, which recently underwent a transition between operators. The Brown family is looking at their legal options.


This tale is fictional—but it is a possibility, however remote. Consider this, though: That possibility becomes less remote when software inspection and data profiling is cut short.

We all live with software bugs and bad data. Our TV systems freeze; our alarm services go off for no reason. PC disks crash, tablets stop working, and cell phones go up in flames. There is no such thing as the perfect software, database, or system—but when the application is mission-critical, comprehensive software inspection and serious database profiling efforts are imperative.


The NPAC system is one of those applications. It is a platform which contains the authoritative routing instructions for 750 million U.S. telephone numbers, and is updated over a million times every day when consumers and businesses change telephone companies, new telephone numbers are added, and service provider networks are updated. The Standish Group estimates that implementing a new system the size and scope of the new NPAC system and database will create 230,000 original defects. The chart below shows how these defects can be found and corrected by comprehensive quality control method.

Testing/Profiling	Normal	Defects Found over Time		
		Long	Average	Short
Step 1: Unit/Column	230,000	207,000	184,000	161,000
Step 2: System/Dependency		20,700	36,800	48,300
Step 3: UAT/Redundancy		2,070	7,360	14,490
Production (6 months)		207	1,472	4,347
Latent defects		23	368	1,863

Neustar to iconectiv



iconectiv to Neustar



So, the solution is obvious: Comprehensive quality control on a new system will dramatically reduce defects. Alternative methods, such as trickle-down migration and conversion can be safer, simpler, less time-consuming and less costly control method - these methods have been successfully used in the stock exchange industry as it moves from online trading systems to faster execution systems. (This is well depicted in Michael L. Lewis's *Flash Boys: A Wall Street Revolt*.)

Under trickle-down migration Sally Brown would still be alive, once the new NPAC was ready to be tested and deployed, it would go live in parallel with the existing on-line system. The older system would process transactions as it normally does. Then, using middleware services like IBM MQSeries, it would capture

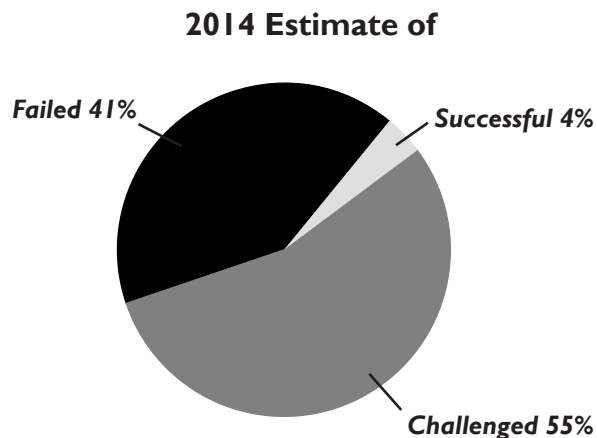
the native transaction and convert it to the format of the new system, which would execute the transaction as if it were live. As defects were found in the new platform, they would be corrected, along with any performance issues – but without impacting users like Sally. Once the new system is deemed to be ready for production, MQSeries would then be turned around, and the old system would behave like a Live-Live back-up—in complete synchronization (1).

Unfortunately, the new NPAC will not be implemented using trickle-down migration and conversion – decision-makers have opted instead to go with an approach known as Big Bang, where the new system is turned up and the old system is shut down, with no parallel processing at all.

Abstract

This paper is written in several parts. After the dramatic opening, a simple solution to the problem is introduced, followed by a condensed overview. An extended explanation (denoted by EE#) will follow, accompanied by a list of references.

The paper is also written in accordance with the academic case study method (EE1)—beginning with an overview of the current NPAC migration and conversion project, followed by an extended explanation and an exploration of the background, the current state, and the future of the project.



New NPAC Project Resolution

The original estimate of the chances of success for the NPAC Project, based on 100 similar projects in the CHAOS Database. The status of the NPAC project is challenged, because it is not considered to be proceeding according to the planned timeline. Note: Traditional Resolution (OnTime, OnBudget, and OnTarget).

In this paper, we will consider factors that may prevent a resolution of the NPAC transition. These factors include Decision

Latency Theory, Project Sponsorship, Database Migration, Testing, and Quality. We conclude with a set of recommendations to improve the project's odds of being completed without additional challenges.

A central theme of Jim Johnson's book *The Public Execution of Miss Scarlet* is the sudden downward trajectory of a large project—in other words, the situation in which a process goes smoothly until there is a sudden “falling apart,” and project members begin to abandon ship. We want to understand why this is so common, and whether upon close examination the NPAC transition project shows the telltale signs. We also want to compare the NPAC transition against the findings summed up in CHAOS Report 2016: “The Winning Hand,” which explored the factors most involved in the success of a software projects. Lastly, we want to consider remedies to ensure—or at least make more likely—that a project like the NPAC transition is successful.

Background

In 2013, in the shadow of the Obamacare website failure, The Standish Group started look at another large and looming project. This project was the possible transition of the Number Portability Administration Center (NPAC) from the current vendor, Neustar, to a new vendor, iconectiv. NPAC is the US's authoritative routing platform for telephone number-based communications, and provides the ability for users to keep their telephone numbers while moving to another telecommunications provider (EE2).

For the last 19 years, Neustar has acted as the operator for NPAC, custom-developing the application as well as the database (EE3).

Current State of the NPAC transition

Publicly available reports and webinars offered by PwC, the industry's Transition Oversight Manager (TOM), paint a muddled picture of the transition's overall status (EE4).


Future of NPAC

In Jim Johnson's book, *The Public Execution of Miss Scarlet*, the title character is a project manager who covers up problems with the project [2]. This is not uncommon. Jon M. Quigley, a professional project management instructor, suggests that although most project teams are quite capable of informing a project manager of potential and even imminent risks, many project managers do not want to acknowledge the presence of such risks [3] (EE5).

Decision Latency Theory

The Standish Group has determined that the root cause of software project failures is decision latency, as noted in a study presented to the International Institute of Informatics and Systemics, July 11, 2017, Orlando, FL [4]. Decision Latency Theory states, "The value of the interval is higher than the quality of the decision" (EE6).

Effect of Decision Latency Skills





Compares projects with highly skilled decision latency teams against projects with poor decision latency teams. Modern Resolution of all projects from FY2008–2017 within the CHAOS database.

Project Sponsor

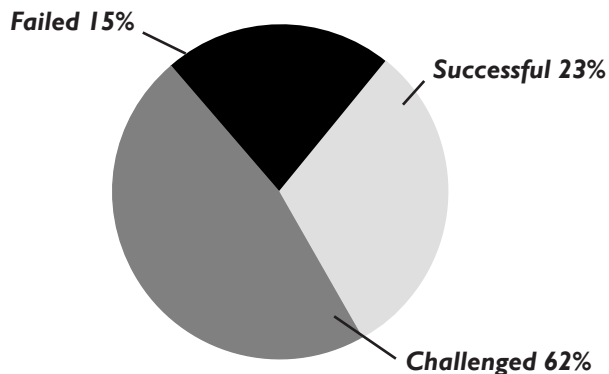
“I think having multiple owners is worse than having no executive sponsor,” says Bill Niemi, former CTO for Fidelity Investments (5). A very critical red flag is having a committee acting as project sponsor and not having a single responsible owner. The Standish Group has for many years held up the quality of the executive sponsor to be paramount in the success of a project (EE7).

Effect of Good Sponsor



Compares projects with highly skilled project sponsors against projects with poorly skilled projects. Modern Resolution of all projects from FY2008–2017 within the CHAOS database.

Data Migration Project Resolution



Modern Resolution (on time, on budget, with a satisfactory result) of data migration projects from FY2008–2017 within the CHAOS database.

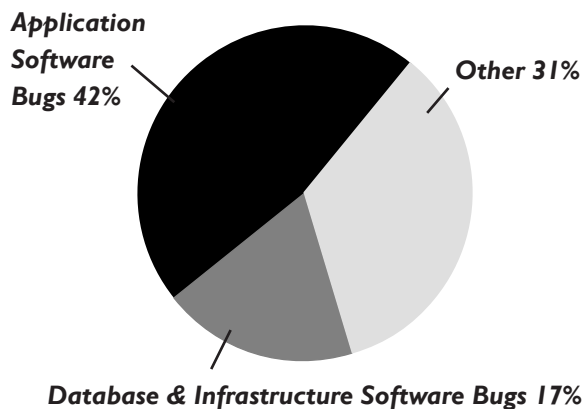
Database Migration

Bill Heil, former president of WebEx Communications, suggests that data is an area that requires a special focus, noting, “Many of the projects that I have seen that have gone bad think about ... the data too late in the project” [6]. Several factors go into a successful data migration project; the Standish Group’s “Migrate Headache” research paper outlines a series of steps needed to ensure success [7] (EE8).

Testing and Quality

One of the clearest signs that a project is heading for the rocks is the truncation of its testing and quality assurance period. This is a critical red flag. The original plan for NPAC was that the testing period, which would include User Acceptance Testing (UAT), would last one year from the final code drop. This period has now been compressed to just two months without UAT. (EE9)

Causes of Application Downtime



This chart shows application bugs to be the leading cause of downtime, as reported in The Standish Group's Rule of Ten research note 2014 (8).

Conclusions

As noted above, the NPAC transition seems from the hazy outside to be going smoothly, but that is no guarantee of future success. We would recommend the following enhancements to improve the chances of a smooth migration:

- The industry should seriously consider deploying the trickle-down migration and conversion method. This would almost guarantee a smooth transition and conversion.
- Project teams should be aware of the detrimental effects of decision latency, and should try to improve decision times.

- Project managers should take a closer look at the condition of the database's third normal form.
- Organizations should appoint a single project owner or sponsor that has experience with this type of application and service. (Following the advice of Former Sprint CIO Dick Lefave, who participated in NPAC's initial start-up, would be a good choice.) This would also help in reducing decision latency.
- A solid roll-back or back-up is necessary. The current state of technology would allow organizations—in this case Neustar and iconectiv—to run in complete synchronization for as long as needed for the industry to be comfortable.

Extended Explanation

EE1: Case study methods are designed for distinctive situations in which there are many more variables of interest than data points, and as a result, case studies rely on multiple sources of evidence in which data needs to converge in a triangulating fashion. This is why case studies typically utilize multiple means of data collection (documentation, observations, interviews and/or secondary data). Case studies are used both for confirmatory purposes (theory testing) and for exploratory purposes (theory building). (9)

This case study research has a confirmatory purpose—to test our hypotheses regarding factors involved in the success of software projects, such as Decision Latency—and follows a mixed methodology, namely a combination of both qualitative and quantitative methods for data collection and analysis. As states, a mixed “numbers and words” method research is “to encourage stronger inferences, to provide a greater diversity of divergent views and to enable researchers to

simultaneously answer confirmatory and exploratory questions, verifying and generating theory at the same time.” (9)

EE2: NPAC is the US's authoritative routing platform for telephone number-based communications. It provides the ability for users to keep their telephone numbers while moving to another telecommunications provider. It connects the networks of over 2000 wireline, wireless, and IP service providers, and delivers updates to its database of 750 million phone numbers to operators nationwide approximately every seven seconds.

Our research was based on the premise that issues and failures in an IT transition of this breadth and magnitude would have ripple effects for market competition, network stability, and public safety. This research later became the basis for a white paper titled “Big Bang Boom” [10]. As a follow-up to “Big Bang Boom,” The Standish Group created a podcast entitled “Why Large Projects Fail,” to be aired on the “CHAOS Tuesday” radio series. We are now following up that project and paper with an interim look at the current state of the Number Portability Administration Center (NPAC) transition project and its progress to date. “Big Bang Boom” and the “Why Large Projects Fail” podcast were an attempt to better understand the root causes of the failure of large software projects.

EE3: For the last 19 years, Neustar has acted as the operator for NPAC, custom-developing the application as well as the database. During this time, NPAC went through the normal growing pains of any software service. By the early 2000s, both application and database had arrived at a mature and stable point, supporting annual growth in data volume and throughput needs without deterioration in service levels. In 2010, the communications

industry and the FCC initiated an RFP process, which iconectiv won by submitting a materially lower bid to replicate Neustar's platform). Neustar accepted the loss of the contract and focused its attention on other areas of its business. Meanwhile, with support from large-scale data migration experts, Neustar created an Office of Transition to support the move.

EE4: It should be noted at the start that from the time the NPAC RFP process began up until now, expectations and processes around the transition have failed to provide a reasonable level of transparency. Most documentation regarding iconectiv's original transition plan, as well as information about its current state of readiness, has either been rendered unavailable or has been heavily redacted. Publicly available reports and webinars offered by PwC, the industry's Transition Oversight Manager (TOM), paint a muddled picture of the transition's overall status.

iconectiv delivered a preliminary version of their new system to selected industry constituents in May 2017, with a second on November 15, and a third and presumably final slated to be delivered on December 15, 2017. Assuming everything goes as smoothly as represented, this last version is meant to represent a complete working model for NPAC. Testing, quality assurance, and user acceptance will take place from December 15, 2017, through February 28, 2018, with the first production launch scheduled for April 8, 2018. The industry's original plan for transition allowed for a full year of user testing prior to launch. The current schedule, however, allows for only two months, and takes place during the holiday season. TOM has reported that this schedule is achievable, but only in the event that no material issues are discovered with the system.

EE5: This situation is not at all unprecedented. In fact, it is very common. Here is some example of a few projects that looked like they were progressing nicely—until they failed. These include the state of Pennsylvania’s \$110 million IT project to upgrade its unemployment compensation system (11). Three projects from the Commonwealth of Massachusetts—Quest, DOR Masstax2, and RMV (ALARS) Modernization—were also failures (12). Two projects from the City of New York (NYCAP and CityTime) failed (13), as did the £46m i6 Scottish police system (14).

It is possible that the executive teams responsible for these projects were so focused on their progress that they did not take into account the results of problems such as decision latency, poor sponsor support, or the quality of data migration and testing. However, these are the very factors that can sink a project faster than missing a major milestone.

Note that this is not to say that the current NPAC project, as depicted in the “Big Bang Boom” paper, is in trouble, nor does it mean that NPAC’s Transition Oversight Managers are wearing rose-colored glasses. It does mean that it is interesting to consider the future state of the project in the light of new research from The Standish Group and other research players.

EE6: Decision Latency Theory states, “The value of the interval is higher than the quality of the decision.” In other words, project teams and project management exhibiting short latency have a much higher rate of success than do teams and management exhibiting long latency periods. The Standish Group has determined that chances of success or failure may increase or decrease by as much as 25% based on an organization’s decision latency skills. There does seem to be

some indication that the NPAC project may be suffering from poor decision latency. One of the things we know for sure, as the project gets closer to being turned over and Neustar is turned off, is that the need for rapid decisions will become more important, not less.

EE7: The Standish Group has for many years held that the quality of the executive sponsor is paramount in the success of a project. A good executive sponsor helps reduce decision latency, and thus the chance of project failure. Many private and public organizations have made having a single responsible owner a requirement for the approval of a software project. There is some consensus around the skills and temperament needed to be a good project sponsor. (These skills are outlined in Jim Johnson’s book *The Good Sponsor* [15], and can be measured through the assessment that is featured in the book.) Moreover, educating a project’s sponsor has been shown to improve the likelihood of its success by 15% to 20%. Dick Lefave, former Sprint CIO, suggests that an adequate investment has to be made along with that sponsorship. In other words, a business owner or executive sponsor has to “stay with it” and shepherd the project through to its successful end (16).

EE8: NPAC’s success will be measured by the successful migration of the old database to new systems. This success often depends on the quality of the database schema to its third normal form. During a recent visit to the Massachusetts General Hospital, we found that the data migration from their existing databases to the new EPIC system had experienced major anomalies. Some of the databases migrated without too many issues, but for other databases, were not able to be converted, and others experienced major errors that made the information unreliable.

In 2015, United Airlines experienced a major system outage that caused flight delays and cancellations. This was the direct result of a problem caused by migrating and merging two airline databases (17). Applications and databases over time decay through normal maintenance and changes, and since NPAC's application and database span the course of 17 years, it would be unusual if that database schema has remained in perfect third normal form. Michael Stonebraker, Dong Deng, and Michael Brodie explain this in their paper, "Database Decay and How to Avoid it." (18)

Our data shows that 3% of downtime is caused by data errors, but it is much higher as a percent of catastrophic failures. This is because the relationships between data elements are not always obvious and are subject to misinterpretation. The only remedy is intense data profiling, laborious manual inspection, and long testing periods. The Standish Group estimates the new NPAC will suffer from as many as 30,000 bugs.

The Standish Group has concluded that a mission-critical relational database in migration and conversion will have an average error rate of 2.2% of errors per row. That translates to about 50 million errors for the 2.2 billion-row NPAC Database. Many of these errors will be minor, such as misspelled names or variables, example (I.B.M., IBM, IBM Corp.) and incorrect addresses. However, errors in the key will account for about .2% during a migration and conversion. Key errors for the NPAC project will translate into roughly 2 million erroneous telephone numbers.

The most important potential problems involve conditional errors. Depending on the state of the schema, these range from .0001% to .0005%—or 200 thousand to one million for the NPAC database. However, even a

very clean schema will still have thousands of conditional errors. These types of errors can surface right away or lie in wait for years. While they are rare, these are the types of errors that caused the United/Continental outage.

The research is based on three sources. The first is a database of 2,000 systems by types of outages, from 1995 to 2012. The second is a profiling study done by The Standish Group in 1999, which looked at 100 data migration projects. The study found that an error rate of 26% by attribute or column. The third source involves studies from the Data Entry Management Association (DEMA) that show the average error rate for unverified data entry is around 7% per record (verified data came in at about 4% per record).

EE9: It is currently very popular for new applications to be released and tested by users, and this is especially true for mobile or consumer applications. However, applications like NPAC, which is essentially a foundation for the US telecommunications system, are a different situation altogether. Consider that bugs cause downtime. Bugs can happen in even old code that has run for years—such as the 20-year-old application bug that brought down the New York Stock Exchange in 2015, causing \$64 billion in damages. New code will have many more bugs and will need time to mature.

A database column error coupled with a software bug can cause a cascading failure that can crash an entire telecommunications system. This is not hyperbole or being over-dramatic; such things have happened. In addition, hackers can bring down a system by exploiting a bug, or even a feature. This is why the information and communications industry's best practice has been to run parallel systems for months, if not years.

REFERENCES

- [1] **Continuous Processing Architecture Presentation** by James Johnson at High Performance Transaction Systems Conference
http://www.standishgroup.com/sample_research_files/HPTS-CPA.pdf
- [2] **The Public Execution of Miss Scarlet**, James Johnson, 2006
<http://www.lulu.com/shop/jamesjohnson/the-public-execution-of-miss-scarlet/paperback/product-22576678.html>
- [3]* **Rose-colored Glasses**, Jon M Quigley, 2014
<http://www.pm2go.com/index.php?r=advice/view&id=414>
- [4] **CHAOS Update, Nanocourses, the Winning Hand, and the Root Cause of IT Project Failure**, James Johnson, 2017
<http://www.iiis.org/Videos2017b.asp?season=summer>
- [5]* **Multiple Owners** Bill Niemi 2014
<http://www.pm2go.com/index.php?r=advice/view&id=367>
- [6]* **Data Early and Often**, Bill Heil 2012
<http://www.pm2go.com/index.php?r=advice/view&id=1167>
- [7] **Migrate Headache**, The Standish Group, 1999
http://www.standishgroup.com/sample_research_files/migrateNew.pdf
- [8] **Rule of Ten**, Standish 2014
http://www.standishgroup.com/sample_research_files/RuleTen.pdf
- [9] **Scientific Research in Information Systems**, Recker, J., Springer Heidelberg, 2013
http://www.standishgroup.com/sample_research_files/Methodology.pdf
- [10] **Big Bang Boom**, Standish Group, 2014.
http://www.standishgroup.com/sample_research_files/BigBangBoom.pdf
- [11] **Computerworld**, Patrick Thibodeau
<https://www.computerworld.com/article/3180325/it-industry/pennsylvania-sues-ibm-over-troubled-110m-it-upgrade.html>
- [12] **Massachusetts Projects**, Senate
http://www.standishgroup.com/sample_research_files/MassProjects.pdf
- [13] **CHAOS Manifesto 2013** Standish
http://www.standishgroup.com/sample_research_files/CM2013-8+9.pdf
- [14] **Government Computing**, David Bicknell
<http://devolved.governmentcomputing.com/news/waterfall-approach-damns-46m-i6-scottish-police-system-5758930>
- [15] **The Good Sponsor**, James Johnson 2016
<https://www.standishgroup.com/store/summer-book-sale-p25.html>
- [16]* **Business Ownership**, Dick Lefave 2014
<http://www.pm2go.com/index.php?r=advice/view&id=960>
- [17] **Reuters** Jeffrey Dastin
<https://www.reuters.com/article/us-ual-flights/computer-glitch-halts-united-airlines-flights-for-two-hours-idUSKCN0PI1IX20150708>
- [18] **Database Decay and How to Avoid It**, Brodie & Stonebraker 2017
http://www.standishgroup.com/sample_research_files/dbdesign.pdf

* Note: Access requires you to register and login (no cost)