

The Standard Autonomous File Server, A Customized, Off-the-Shelf Success Story

Susan K. Semancik and Annette M. Conger

NASA Goddard Space Flight Center's Wallops Flight Facility, Wallops Island, Virginia,
USA

{Susan.K.Semancik.1, Annette.M.Conger.1}@gsfc.nasa.gov
<http://www.wff.nasa.gov/~websafs/>

Abstract. The Standard Autonomous File Server (SAFS), which includes both off-the-shelf hardware and software, uses an improved automated file transfer process to provide a quicker, more reliable, prioritized file distribution for customers of near real-time data without interfering with the assets involved in the acquisition and processing of the data. It operates as a stand-alone solution, monitoring itself, and providing an automated fail-over process to enhance reliability.

This paper describes the unique problems and lessons learned both during the COTS selection and integration into SAFS, and the system's first year of operation in support of NASA's satellite ground network.

COTS was the key factor in allowing the two-person development team to deploy systems in less than a year, meeting the required launch schedule. The SAFS system has been so successful; it is becoming a NASA standard resource, leading to its nomination for NASA's Software of the Year Award in 1999.

INTRODUCTION

Deciding to use a commercial off-the-shelf (COTS) product as the basis or cornerstone of your system software design is a risky business. Among the strongest fears is the "unknown" component either of the product or the vendor. High among concerns using COTS products are the following:

- What if the vendor goes out of business or drops the product you've chosen?
- What if future versions of the product change or eliminate features you were depending on or around which you built your application?
- What if the product does not operate/function as advertised (and you don't discover this until you are deep into your development/schedule)?
- What if the product has errors/bugs that the vendor won't/can't correct, or is willing to correct, but not in time to meet your schedule?
- What if future versions won't operate on your platform, or version of the operating system, or become incompatible with your hardware components or drivers? (And these new versions contain bug fixes or features you need?)

All of these questions are even more critical if you are considering a new COTS product or version and you perform all of your initial investigation with a demo or

pre-release version of the product. Just replace the words “future version/product” with “the new release” for equally troubling COTS concerns.

PROJECT

This was the situation in which we found ourselves in the summer of 1997, when the assignment was given to design, develop, deploy, and field-test an autonomous file server (Standard Autonomous File Server - SAFS) at National Aeronautics and Space Administration (NASA) ground stations in time to support the Quick Scatterometer (QuikSCAT) satellite launch planned in less than one year. This system had to manage distribution of satellite files to customers of near real-time data without interfering with the assets involved in the acquisition of the data at the ground stations, and the processing of the data by the customers. By the Fall of 1997, the SAFS team of two had budgeted the timeline as shown in Figure 1.

This gave us roughly five months for design and prototyping, three months for procurement and development, and three months for shipping and personally installing systems at NASA Goddard Space Flight Center in Greenbelt, Maryland; NASA Wallops Flight Facility (WFF) at Wallops Island, Virginia; Poker Flat Research Range in Fairbanks, Alaska; and the new satellite tracking station in Svalbard, Norway.

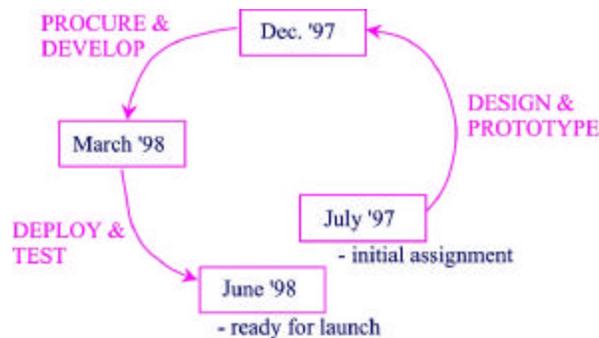


Fig. 1. SAFS Project Timeline for QuikSCAT

DESIGN

We designed a system that would allow the SAFS to operate as a stand-alone solution, monitoring itself, and providing an automated fail-over process for enhanced reliability. Soon after the initial assignment, we began a search and evaluation process for COTS products, not only to help meet the schedule, but also because it is NASA's policy to use COTS software and hardware products wherever possible to save time and money, as well as to re-use government-developed products in the most

efficient manner. With our aggressive schedule, we desperately needed a COTS software product that would provide the reliable, guaranteed file delivery part of our design, and COTS hardware for speed, reliability, and redundant, hot-swappable storage.

We were able to pattern part of our software design for automated file handling and messaging on a system developed at WFF, and concentrated our COTS software search for a product that would provide a quicker, more reliable, prioritized file distribution.

We created a prioritized list of features desired in the COTS software product to help us better evaluate available products:

- Reliable, guaranteed file delivery
- Recovery from point of failure
- Multi-platform support
- Stop-resume transmission control
- Auto-detection of incoming files
- Processing flexibility:
 - Multiple distribution points
 - Pre-/post- processing capability
 - Alternative actions on failed transfers
- File transfer security
- Programmable bandwidth

COTS RESEARCH

We searched the Internet and talked with local experts having experience in similar areas. But while the Internet was helpful, it did not yield a comprehensive list. During peer and preliminary design reviews, especially those including people from other NASA Centers, we gained additional insights and sources to consider. This is actually how we were pointed in the direction of the product we eventually chose.

The Internet is very useful in gaining detailed information about the products and vendors under consideration during a COTS search, and in some cases, even getting demonstration versions of products. It is imperative to obtain demo versions of the COTS software or loaner COTS hardware whenever possible to be sure features are as advertised and that the learning curve to use the product will meet your timeline. It is also important to get references from vendors of the customers that are using their hardware or software in similar situations. By contacting these sources, you may be able to find out if they had product or vendor problems and if so, how easily were they resolved, and if there are any configuration/use limitations with the product before actually committing any of your resources.

COTS SOFTWARE

At the end of our COTS software product investigation, we had a list of several products with features similar to those contained in our list. Some had only a few of our desired features, and others had more capabilities, such as multicasting, which were not part of our project's requirements. A "lesson learned" that worked very much in our favor was to select a product appropriately sized for our application, and a vendor whose size did not inhibit a working relationship with us. This way we did not pay for more features than we needed. In general, a product with a smaller feature set is more likely to be less complex for the vendor to maintain, thus giving faster responses to bug fixes. Also, if the product is closely aligned with your project's requirements, the enhancements you suggest may fit into the vendor's development plan, also resulting in faster upgrades.

Our development process started with a demo copy of the COTS software product that best matched our requirements. We followed on-line vendor tutorials and documentation to gain a good foundation in the use of the product. By this simple approach (which some developers skip to save time, but which usually results in lost time due to false starts and lack of overall understanding of the product's capabilities), we gained insight about how to best incorporate the product into our design. During this initial period, we developed a working relationship with the vendor as he responded to our inquiries for clarifications about more complex procedures, especially involving fail-over strategies. The vendor's willingness to assist us and to extend our trial period while we determined how well the product fit with our requirements, were both good indicators of the level of support we could expect after purchasing the product.

During this trial period, we were able to demonstrate the ease of integration of the COTS software product with the re-used software scripts and the COTS software processing flexibility. We purchased different hardware versions of the COTS software product in order to simulate in a lab how the product would work on various customer's platforms, as well as to model the field operational environment. What we learned from this prototyping was passed on to project customers, helping them to reduce their learning curve with the product, and making them a more willing partner in our development effort. We also encouraged them to purchase vendor support as we did, because of the time and sanity it could save.

To accommodate those customers who decide not to use this COTS software for file transfer with the SAFS system, we built in an alternative option to use File Transfer Protocol (FTP) for their file acquisition from the SAFS. Most of our projects' customers choose to use the COTS software because of the added security, reliability, and guaranteed delivery that it provides through our system.

COTS HARDWARE

The COTS hardware search concentrated on servers and redundant array of independent disks (RAID) components that would be robust, reliable, and expandable, meet our speed requirement, and be maintainable in remote locations. At the time, our investigations found the fastest server and RAID drive systems were not available from the same vendor. This led us to a dilemma: we could either get all components from the same vendor and not meet all our performance requirements, or get the best components from multiple vendors, and possibly have configuration or compatibility problems later. After much discussion in peer and pre-design reviews, we were able to convince any opponents that the latter decision was the best. While this approach is likely to be more expensive, it gave the results we needed to meet our aggressive schedule. We did have a few instances in our early development where it was more difficult to track hardware/configuration problems to the specific component because of the multiple vendors, but the performance aspects of the system far out-weighed this difficulty.

The expansion capability of any COTS hardware system and the level of vendor support needed, both during development and deployment, can heavily influence the product you choose. The RAID vendor we selected had the fastest and easiest system to expand, and also used an external personal computer (PC) in their design to free the server from the RAID monitoring and configuration tasks. With their worldwide network of support personnel, they could provide a field engineer to accompany us during field installations to optimally configure our systems, which greatly helped us in remote locations such as the satellite tracking station in Svalbard, Norway. The integration of COTS software with the configuration of COTS hardware from different vendors can be a significant effort. Whenever possible, it is important for the design/development team to personally perform on-site installation of their systems. It gives the team concrete knowledge about field configuration of the system, an appreciation for the operating environment, and an opportunity to develop a rapport with the staff for future problem resolution.

VENDOR SUPPORT

We found it was crucial to have support/maintenance contracts from both COTS hardware and software vendors through our development, deployment, and first year of operation in order to have quick resolution to problems, and to assist in optimally integrating and configuring the systems. The first year of operation is normally a “shake-down” period that tests your system to the limit in situations not always possible to predict or duplicate in a prototyping environment. Under normal conditions, no operator involvement is needed for the SAFS. Our biggest problem during operations came from an unexpected source – the operators of an external system that normally sends data to the SAFS through an automatic process. While we designed the system well to handle automated processing, the ever-changing operational environment at the ground stations (both commercial and NASA) led to occasional operator errors when they needed to perform manual transfers from their

systems to the SAFS. This put us in the mode of training new personnel in correct procedures to follow to avoid problems with options we are using within the automation part of the COTS software product. We developed an early warning system that will alert us when such errors are occurring so we can manually correct them before it causes a system problem while we explore options for an automated solution.

PROTOTYPING

It is important to prototype your system's hardware and software in a lab setting as similar to the field environment as possible. Testing should be ongoing while your design matures in order to improve the design and to identify any problems while you are still in the development stage, rather than in field-testing, at which stage it may be too costly or impossible to retrofit a solution. Utilities created to help validate and verify development efforts should be considered as tools useful for operational assessments as well. For example, while developing in the lab, we created a display that visually indicates the file transfers and message interactions of the systems as they occur. This not only helped us in our development effort, but also was especially helpful in our project readiness demonstration. It was so indispensable, that it eventually led to the creation of an automated web site that continually reports on the operational file transfers and message interactions at both the ground station and customer levels so all users can track the SAFS performance.

DEVELOPMENT

Using an iterative waterfall methodology, we developed the prototype in stages in the lab environment, with ongoing integration and testing through the design's maturity. The initial phase was critical, since we not only learned how to master the intricacies of the products, but also were successful in prototyping a system to handle file transfers for single project support. This was the first version released to support the QuikSCAT project, with Figure 2 illustrating our software configuration.

After QuikSCAT launched, we had to maintain our operational systems in the field while we continued working on phase 2, multiple project support. This involved analysis of both feedback from our end users and expanded requirements to handle more projects desiring to use the SAFS. Additional projects meant new deployments at the satellite tracking station in McMurdo, Antarctica; and at the University of Alaska in Fairbanks, AK; and the need for COTS hardware, operating system, and software upgrades and enhancements. The second phase of our design needed a more robust, generic system, with a customizable priority scheme to handle multiple projects. After discussing the possible techniques for accomplishing this with the COTS software vendor, he enhanced his product by implementing a file priority parameter that would allow files of equal priority to transfer using shared bandwidth, and files of lower priority to suspend transfers until higher priority transfers completed.

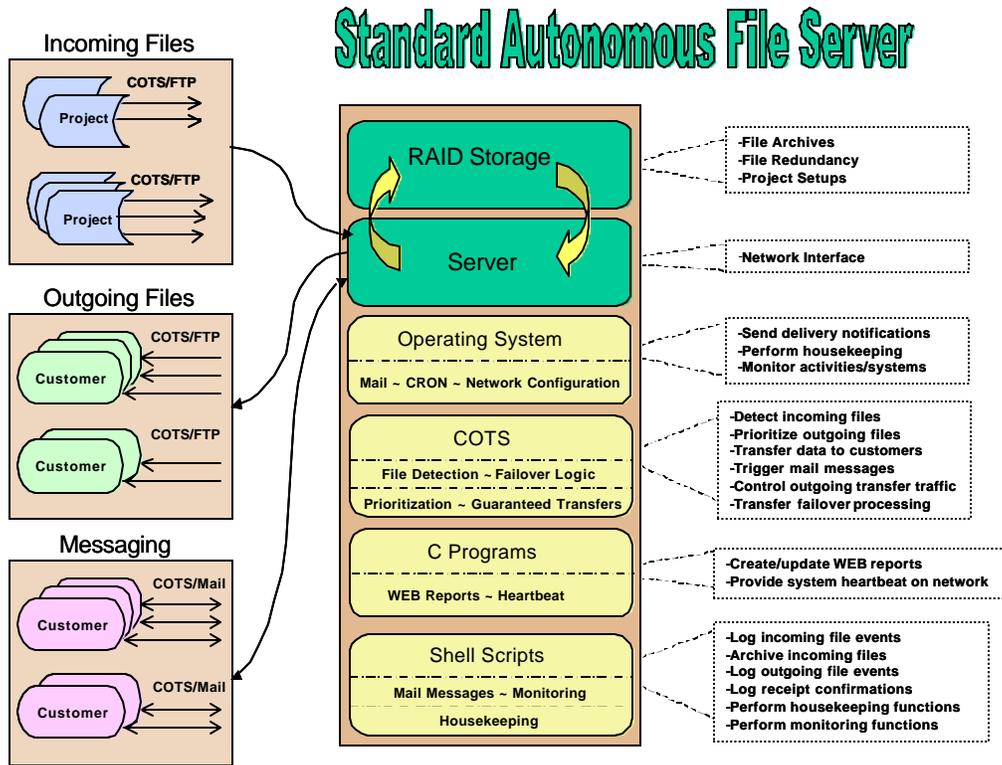


Fig. 2. SAFS Component Design

ENHANCEMENTS

What helped us immensely in being able to handle everything in a timely fashion was keeping a prototype system in our lab. We used it to test enhancements to the system, and to configure the COTS hardware upgrades before field installation. Problems discovered in the lab were easier to resolve because resources were more accessible and operational systems did not have to be disrupted. Both of the COTS hardware vendors had the desirable feature of also being able to remotely access their components for debug/problem resolutions in the field. One desirable feature in a COTS software product is its ability to perform internal logging. Since the internal operations of a COTS software product are often hidden from the user, this feature may be the only way to trace errors or define the point within the product at which they occur, thereby getting speedier resolution to problems during either the development or operational periods. Though this can also be a possible source of

some problems. For instance, one problem did not show up until the systems had been running for about a year. In that time, the quantity of file activity had generated so many logging files that it was causing system errors and poorer performance. It was at this point that we learned there were COTS software housekeeping functions we needed to perform on a regular basis to keep the system operating optimally.

As a system matures and expands, it is important not to approve all requests for additional options by customers or new projects that come on line. We tried not to make concessions that would compromise the performance of the system or would make the design less generic and more difficult to maintain. We did have to make some adjustments to handle project file names as well as the SAFS naming scheme initially developed for QuikSCAT support. These changes were accomplished in a reasonably short period because our design was flexible and modular in nature.

LESSONS LEARNED

Table 1 illustrates the lessons we have learned with the SAFS project and how these lessons impacted our design, development and maintenance efforts.

Table 1. SAFS Lessons Learned

LESSON	IMPACT
Use COTS products and re-use previously developed internal products.	Shortens development time.
Create a prioritized list of desired COTS features.	Focuses the COTS evaluation effort for a better decision.
Talk with local experts having experience in similar areas.	Helps to identify additional resources to explore or re-use; improves the design.
Conduct frequent peer and design reviews.	Improves the design; provides early identification of changes to either the project requirements or operational environment.
Obtain demonstration versions of COTS products.	Assures features are as advertised; determines if product learning curve will fit into project timeline; helps identify configuration/use limitations.
Obtain customer references from vendors.	Helps identify previous product, configuration, or vendor problems, and how easily they were resolved
Select a product appropriately sized for your application.	Reduces cost, design complexity, and maintenance of product.
Choose a product closely aligned with your project's requirements.	Results in vendor being more likely to incorporate requested enhancements into the product, resulting in faster upgrades.

Select a vendor whose size will permit a working relationship.	Improves vendor response time for requests for clarifications and help with advanced applications.
Use vendor tutorials, documentation, and vendor contacts during COTS evaluation period.	Results in time saved by gaining insight into how best to incorporate the product into your design; provides baseline for level of vendor support to expect after purchase.
Prototype your systems hardware and software in a lab setting as similar to the field environment as possible; simulate how the product will work on various customer platforms; model the field operations; develop in stages with ongoing integration and testing	Helps to identify problems while still in the development stage, not in operations when it may be disruptive or not possible to retrofit a solution; provides a more mature design resulting in fewer problems in the field.
Pass pertinent information on to your customers	Helps to reduce their learning curve with the COTS product, and makes them a more willing partner in the development effort.
Accommodate your customers, where possible, by building in alternative options	Provides flexibility and modularity to the design, making the system more robust and generic.
Don't approve all requests for additional options by customers or new projects that come on line.	Avoids compromising the performance of your system, making it less generic or more difficult to maintain.
Select the best COTS components for product performance even if they are from multiple vendors.	Promotes product performance in lieu of design simplicity.
Consider the expansion capability of any COTS product	Ensures ease of future integration without redesign.
Determine if the vendors support is adequate for your requirements	Enables worldwide on-site personnel for hardware support in remote sites.
Personally perform on-site installations whenever possible.	Gives the team concrete knowledge about the system's field configuration, an appreciation for the operating environment, and an opportunity to develop a rapport with the field staff.
Have support/maintenance contracts for hardware and software through development, deployment, and first year of operation	Saves time and your sanity; use support to optimally configure and integrate the COTS product into your system
Create visual representations of system interactions where possible.	Helps during development effort, in demonstrations of the project's readiness, and provides a prototype for an operational utility.
Obtain feedback from end users	Helps to identify problems early;

	provides a more flexible design; gives an indication of system performance during operations.
Maintain the prototype system after deployment.	Provides a non-operational test-bed for enhancements and for configuring upgrades
Select COTS products with the ability to do internal logging	Helps trace errors or define the point within the COTS product in which they occur; produces speedier resolution to problems.

SUMMARY

In summary, we were able to mitigate some of the risks/concerns we had with using COTS products by considering a vendor's history and reputation through their customer's feedback, our success with trial versions on multiple platforms, vendor support during evaluation periods, modular design, prototyping, maintenance and support contracts, frequent contacts with vendors and customers, peer and design reviews, constant testing, on-site spares for operational backups, and re-use of successful operational software. We found the addition of an operation's contractor assigned to operational system administration and maintenance responsibility helped greatly in allowing the team to complete development of phase 2. For those desiring more details about the SAFS design, development, and deployment phases, please see the paper at the following link: www.wff.nasa.gov/~websafs/iafpaper.pdf.

The SAFS project was successfully transitioned to an operations contract this year, January 2001, which was made easier by the success and reliability already proven by the SAFS system in support of QuikSCAT and Earth Observing-1 (EO-1) satellite missions. Our successful integration of COTS products into the SAFS system has been key to its becoming accepted as a NASA standard resource for file distribution, and leading to its nomination for NASA's Software of the Year Award in 1999.
