

New Application: What is the Network Impact

Robert L. Godfrey, Jr
NDIC Fellow
rlgodfr@dia.mil

Abstract

Do you know what the impact of your Service Oriented Architecture (SOA) deployment is to the operational network?

During deployment, the network requirements for the application are discovered. Deploying functions directly to the operational network forces the network technicians to quickly adapt the network to these requirements. Since this is not optimal, we need an improved process.

A way of improving this process is to use a test network that simulates the operational network although a better solution would be to extract network requirements and verify the requirements using the test network. The test network reduces operational impact by removing the development of the requirement from the operational network. To enhance this process farther would require that the network requirements be extracted during development. This process would use a common language between the developers and network technicians that capture the network requirement. The test network would then be used to verify the requirement of the new function before it is deployed to the operational network and reduce the impact to operations.

1. Introduction

The Department of Defense is looking to Service-Oriented Architecture (SOA) to deliver reusable mission and business functionality as standardized blocks [1]. SOA allows developers to rapidly develop functions and provide that functionality to any application that comes aware of the function [2]. Development methods do not consider the impact a new function will have on the operational network, when it is deployed. The network impact needs to be considered before new functionality is incorporated into the network. This impact needs to be considered as part of the development lifecycle. The Department of Defense Architecture Framework version 1.5 has

added fields for SOA network requirements. This paper is aimed at the deployment process and only advocates finding the necessary data in the development versus discovering the data during deployment.

The current deployment process can be described as using our operational network as a testing center for incorporating new functionality. The deployment process does not include any network impact testing and forces the network technicians to quickly modify the network to support the new function or to re-enable an older function that breaks. This process has placed a burden on network staffs to maintain the current systems. The network impact must be quickly resolved to restore any impacted operational systems and to provide reasonable performance for the new function.

Since direct deployment to the operational network is not the most optimal way to deploy new functions, then we need to look for an improved process. Two alternatives come to mind. The first is to use a test network that simulates the operational network and the second is to develop a methodology that extracts information about the functions network requirements to develop a deployment plan.

The test network improves on directly installing the function to the operational network for testing that impacts users and operational systems. The testing network still utilizes the same process but on a separate network from the operational systems. This still relies on running the application and blindly making network changes based on observations. A better process would be to extract the network requirements from the function to use in verifying the network capability.

The extraction process would require using a common language for defining a set of parameters. Developers and network technicians would have a common understanding of the requirements and how the function is expected to operate in the network. This common language would improve the process by not impacting the operational network to test a new application. The testing would not be random modifying of the operational network until you find a

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 21 MAY 2008		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE New Application: What is the Network Impact				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) National Defense Intelligence College				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES AFCEA-GMU C4I Center Symposium "Critical Issues In C4I" 20-21 May 2008, George Mason University, Fairfax, Virginia Campus					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

solution to the performance issues. Also, the test network would enhance the solution because it would allow focused testing with pre-knowledge to what is needed without having to install the software, record the results, and then find a solution to any performance issues.

The network requirements would be known and the technician would know where the network is not capable of supporting the new function and can work toward finding a solution. After testing the results in the test environment, the transition to the operational network will be documented and verified providing a smoother and faster transition of the new function. The network impact is not new to SOA but to all types of developments.

There is still a lot of research that needs accomplishing to fully develop this extraction process. The results of this process have not been verified and tested against a real development. The network is not fully understood and processes would have to be developed to ease the transition to extraction methodology to replace the trial and error method.

A lot of discussion and work has gone into improving how we build applications but little has been spent on deployment. We started with simple monolith programs, transitioned to client-server, and now are focusing on web deployment. As we progress to a System of Architecture, we need to look even more closely at how this is going to impact our networks.

The purpose of this paper is to raise awareness about our current processes for incorporating new functionality into our networks and outline a new way of approaching the deployment challenge in the development phases. So, are we maximizing our network technician's time by diverting this scarce resource to help in development of new functionality?

We need to improve the current process and begin studying what it would take to improve the deployment cycle and produce information that could allow the networking staff to verify the networks demands of the functionality before it is activated on the network.

2. Current Process

The current process revolves around the development of new applications. The following diagram will serve as a starting point for talking about development networks.




Figure 1

Figure 1 is representation of an operational network with a closed development network. The local area network has servers providing file shares and email to the developers. Most of the development is done within the DEVNET circle containing a set of developer workstations and application servers. The local area network consists of a switch that provide high-speed service between machines and the edge router. The router connects to the wide area network at a much slower speed than the local area network.

2.1. Closed Development Environment

The developers utilize a network that connects development team members. The network is usually a high-speed network, which will be referred to as the DEVNET in this paper.

DEVNET is considered a closed network because the developers use their local machine to develop and test new code. The development cycle follows a series of defining requirements, developing specification for the code, writing the code, and unit/bug testing. This cycle is repeated until the code is ready for acceptance testing. Testers/users are given access to the software on DEVNET to verify if the software performs to the requirement.

The process is accomplished inside the DEVNET and network performance is based on the high-speed of the network, which does not always match the operational network capabilities. Although some organizations have development hardware and separate test hardware, the network is the same network. This use of DEVNET does not change the result.

Because the development is on a 100 megabit connection (figure 1), the testing will reflect that it will work on a 100 megabit end-to-end connection. Which based on figure 1, we can see that a user connecting across the wide area network will be limited to T1

connection at 1.5 megabit. This is almost 1% of the speed that the function was tested against.

2.2. Publish Interface

After passing testing, a new function requires that its interface be published for other functions in a System of Architectures to utilize in building applications. The interface does not provide any network specifications but describes commands and data that the function is designed to use in performing its function.

2.3. Fix Application

Once the interface is published, now applications can be built to take advantage of the new function. The network impact is not available in building an application from various functions. The new function has been tested locally on a high-speed DEVNET but what happens when it changes to the operational network that includes slower wide-area networks?

The function will operate in a very different environment than the environment used by the developers in designing the function. The operational network has many applications vying for a limited amount of resources. The operational network is a more complex ecosystem.

The operational network has connections between locations, which are slower, than the local area network that supported the DEVNET. The network can have many routers between the requestor and function, which are impacted by jitter and packet loss.

It is the network that impacts the performance of a well developed function requiring changes to the network or going back to the developers to enhance the application to overcome short-comings in the network.

This pushes the task to the network technicians to try and improve the network to support that new function with no requirements on what the function needs to perform its task. This carries a risk because the network was modified for the last function and now changing it again may undo the work to make another function work. This can quickly become a self-defeating exercise leading to a network that is sub-optimal but capable of meeting the minimal functional requirement. This reactive response happens because we try to deploy the function quickly and skip steps and take shortcuts that lead to time-consuming problems and troubleshooting later [3].

2.4. Multicast Upgrade

Offutt Air Force Base was chosen as one of the locations to demonstrate Radio Over Internet Protocol

(RoIP). The project was designed to support linking various radios using different technologies together to support remote aircraft contact, expanding coverage of signals to training area, and to allow maintainers direct access to aircraft for troubleshooting in-flight issues. The project required modifying the network infrastructure to support multicast protocol. The plan for deploying the protocol was discussed and plans were made for the role-out. The configuration was not tested before deployment.

As the routers and switches were modified to support the new function, operational systems started to experience problems. Network technicians were forced to rapidly respond to outages without any knowledge on why the network was starting to hiccup.

After spending a large amount of time tracking the issues, the decision was made to restore the old configuration to the routers. This solved the issue but it was unknown why and research was accomplished to track down why the network was sensitive to this change.

The result was to update some equipment on the network to enable the multicast support to be installed. The service was eventually activated to support the RoIP demonstration but not before it caused a loss in productivity.

3. Is there a better way

The current process works but can it be improved to reduce deployment and maintenance time?

The research that I am performing is to show that we can improve upon this practice by adopting a more rigorous methodology. Two possible candidates that could improve the process are to use a test network that mimics the operational network and to extract information from the functions development to use in verifying the operational network.

4. Process changing

The first possible solution to solving our operational network impact is to use a test network to simulate the operational network. Another solution to solving our operational impact is to gain an understanding of the application requirements because without proper requirement gathering we cannot be sure that the network will meet the need [4]. The understanding of the requirements for an application can be used to verify the operational network and would be a good way of bringing the application to the test network. This would reduce the time spent trying to find a problem in testing and to maximize the utilization of the operational network.

5. Test Network

The current process is skewed towards using the operational network for testing a new application. The dedication of a network system for testing would secure the operational network while allowing real testing of the new application. The test network must be built to simulate the network topology and traffic that is present on the operational network.




Figure 2

The test network shown in figure 2 is designed to be an approximation of the operational network. The test network does not need to exactly match the operational network in terms of devices. The goal is to have a network that approximates the physical layout and can provide as realistic of a network clone as possible to the operational network.

5.1. Topology

The test network must start with a good understanding of the current networking environment. This would include having an understanding of the network topology to use in building your simulation. If your application is going to be thoroughly tested, then the path that the data travels in the operational network needs to be represented in the test network. The paths must simulate hardware making up the network to include routers, switches, and hubs. Besides the hardware that makes up the physical network, the network capabilities, including link bandwidth, must be considered to simulate the operational network.

The topology will be the road map for developing the plan. If we don't have the road map of the operational network, then we will not be able to decide the paths that are affected. Therefore, the test network must approximate the paths that traffic will take on the operational network.

5.2. Traffic analysis

The test network's traffic must be designed to mimic the traffic patterns and protocols that are present on the operational network. The traffic does not have to be an exact clone of the operational network but the traffic must approximate the congestion, jitter, and other conditions to allow studying the impacts of the functionality on the network.

5.3. Scenarios

The scenarios to use in the testing phase must simulate the expected uses of the function and resemble the extremes that it will encounter. The stress points caused by the extremes are going to be the factors that we are trying to understand. If the scenarios only test minimal traffic that is expected then the network results that come from the testing will not be useful in tuning/verifying the operational network. The data collected will provide a false sense of the requirements of the application and will not provide the needed feedback to actually verify the network.

The scenarios must also include some tests that show what is expected to happen in the function over time. A specific function may be limited to a small audience initially but can be expected to expand over the next year. That growth needs to be captured and considered in the analysis. The hard question is the load created by growth that you can't predict.

As the function is used by other developers to build their applications, the function owner is unable to predict that traffic. The unanticipated consumer is going to grow the demand for the function and the function developer needs to be able to estimate that behavior to help in predicting the growth. The growth will directly impact the network capabilities that are available. As an application grows, its resource requirements from the network will grow causing the network operators to adjust the network. These changes to the configuration could have an adverse impact on other functions.

6. Application Parameters

The unknown requirements for deployment of a new function need to be turned into known parameters. The developers of a function need to be able to describe their application in a language that can help the network operators verify the operational network is ready for the new function.

6.1. Extract Parameters

A new way of approaching this problem is to use a set of parameters that can be refined in the testing network and used to prepare the operational network. The parameters would describe the application's network resource requirements in a language that is common to network technicians.

The process would consist of three methods to extract the information needed for the testing. Modeling and simulation based on how to optimize a function would provide basic network resource requirements. Concept of Operations (CONOP) that would provide the requirements for what the function is designed to produce and finally the functional requirements that state how the function was designed to handle the data.

The modeling and simulation is going to include calculations that show the bandwidth requirements and how the function will process data. The modeling must include the extremes to verify the requirements. If the modeling is only on an average then the impact to the network will not correctly match the results that will be produced in the production environment. Since the model has to review the complete spectrum for the function then the model must include maximum resource demands. Additionally, results must be found that predict the growth or changes over time. Without the expected changes over time, the function has the possibility to degrade and slowly impact other functions.

The CONOP is going to define the importance of this function and how it is expected to fit in the organization. The CONOP should include a section of the priority of the function and how the function fits into the overall scheme of the organization. The CONOP will also provide the function requirements for performance and help in building the models because it will state the expected demand and growth of the function.

The design to meet the functional requirements is the last method that needs to be included in the process. The design will provide the capabilities of the function to work in extreme cases. The design will implement the fail-over capabilities, the level of network limits that can be overcome, and how the data will be packaged. The design based on the stated functional requirements will set the network limits for the function.

6.2. Possible Parameters

The possible parameters for documenting a function that need to be looked at are Jitter, Bandwidth, Type of

Service, and Class of Service. We need to first define the usage of these terms.

- Jitter – amount of variation in delay that a network introduces [5]
- Bandwidth – measure of the capacity of a transmission system [5]
- Type of Service – amount of variation in the amount of traffic that a function sends in a set amount of time. The traffic will be steady or bursty in nature
- Class of Service – designates the transport network characteristics of a session [6]

The parameters will serve as a common language allowing developers to talk to network technicians about what a function needs in terms of network resources. The parameters will be critical to defining a service level agreement with the user community.

6.3. Methodology

The parameters need to be defined at the beginning of the development process using performance engineering to analyze the expected performance characteristics [7]. The parameters development needs to start with the concept of operation phase of the development cycle. The network requirements need to be identified as early in the development cycle as possible to accomplish the greatest payback. The CONOP will describe how the function will be used and what is the expected growth in the future. Developers can use the CONOP to help in gathering and understanding the functions requirements.

The requirements documentation for the function needs to include a section for network requirements. This will state the expected response times, the priority and class of traffic, and network limitations. From the requirements, the developers will design their code to accomplish these requirements. The requirements will be tested during the test phase to verify the function. The test cases are quantitative based and not qualitative.

The test network will be used to verify that the function can meet the network requirements. If the function needs to have a steady 10 Kilobits per second of bandwidth then the test would verify that the test network provides this capability. The test network can then be modified to verify that the necessary capability is available and this information will be incorporated in to the deployment plan.

The deployment and sustainment plans will include information that the function needs to be deployed to the network and the expected future requirements. Once the function is deployed, the network will need to be monitored and modified based on the sustainment plan.

6.4. Benefits

The benefits from using this type of methodology are defining a service level agreement, quicker deployment of functionality, less interruption to current operations, improved network management, better vendor understanding of needs, and a common language between developers and network technicians.

Service level agreements (SLA) can be based on the results determined in your modeling and simulation and verified in the test network. This gives some confidence in the values and allows a more robust SLA to be created that can reflect reality. Many SLA just include an up time on the network hardware and application providers but does not say much about actually being able to use the function. The agreement can include very specific requirements because we have a higher confidence in the network's capabilities. The SLA will allow professionalizing the network by providing a mechanism to judging the network health.

The parameters will allow quicker deployment because the network requirements for the function will be known and the operational network changes already decided from the test network. The changes can be incorporated in the network, the new function installed, and a quick test to verify that the system is working. The network technician will not have to rediscover the necessary changes in the operational network to install a new function.

Since the network technician will have to worry less about what it will take to incorporate the new function. The newly deployed function will have less of an impact on current operations. The network changes can be planned and implemented in a systematic way. This process will lead to more stability in the network and greater confidence in the network capabilities.

The knowledge of the installed functionality, its impact on the network, and how the network is operating will pay dividends in the management of the network. Process control principle is knowing how your system works by keeping it between tightly controlled tolerances. The knowledge that is accumulated about the various functions using the network will allow the technicians to manage the network more efficiently and effectively. They will have an understanding of the complete system and can route around failures, understand how surges will affect the network, and report on SLA compliance.

If you take the understanding of the network performance and the expected future growth of a function then discussing network enhancements with vendors will be easier. If you can show and explain exactly what is needed to improve or meet new functional requirements then vendor's products can be

evaluated against supporting the documented requirement. Vendors will be able to tailor the product to the need versus trying to discover the requirement, recommend a solution, and then test it.

This common language of talking about requirements for the network will improve communication between vendors and network technicians. The same will apply between developers and network technicians. The ability to use a common language will remove ambiguity and misunderstanding of what exactly is the requirement for the function.

7. Multicast revisited

Lets look at installing multicast at Offutt Air Force Base and how it would have been accomplished using this new methodology. The first thing would have been to identify that multicast support is required to use RoIP. The requirement would have driven the changes to the router configuration files to include the network demands that RoIP would place on network resources.

The router changes would have been installed into our test network environment and the test would have shown that the changes impacted some of the equipment. The operational community would not have noticed the impact because the impact would have been discovered in the test network. The technicians would not have to quickly restore files or expend energy trying to recover operational capabilities.

The technician could have concentrated on finding the reason that the new configuration was impacting the network and finding a solution. After finding the solution, the new configuration and network requirements would be updated in the deployment plan. The plan could then be implemented on the operational network and the network verified to make sure that the changes did not impact any systems.

The use of a common language would have allowed the development of the requirement, captured in the changes to the configuration files, and finally in the deployment plan.

8. DODAF

The Department of Defense Architecture Framework (DODAF) version 1.5 includes some net-centric guidance. The extensions were specifically added to support SOA development.[8] The DODAF is designed to capture the DoD architectures and help describe architecture artifacts across mission operations and processes and move away from a focus on architecture products to a focus on architecture data

[8]. The data is important but only if it is available for those that need it.

The DODAF has several architecture products that include network requirements. The network requirements are listed as part of the SoaServices that is recorded in the AV-2 Dictionary.[9] The AV-2 captures the number of concurrent users, protocol, expected bandwidth, and allowed jitter. [9] This information is available in other views like SV-2 System Communication Description and SV-7 System Performance Parameters Matrix to help in describing the network environment that is needed.

The DODAF is currently not made available to the network technicians who need the information to perform their duties. The information that is provided will serve as a good background for defining a QoS environment needed to incorporate various applications and functions together.

9. Concerns

There are many concerns that need to be examined with this approach to improving how we incorporate new applications. Some of the areas that need to be studied further are: what exactly are the parameters, what are the models needed, and how can this be applied to our current networks?

The parameters that I defined in this paper are what I am proposing to use in this methodology. The parameters were picked based on what the demands of various types of functions have in common. The parameters do not handle some specific function requirements like packet loss, which is important to certain functions. So have all the parameters been identified to properly define the network requirements for a fairly wide range of functions?

The models that are needed have not been designed or tested to verify that they work. Research needs to be done to take our current queuing theory models and expand them to produce results that can be used in defining the parameters. The models have been used to help in optimizing application and networks congestion but have not been applied to finding an applications network requirement.

The current networks have been developed over many years and it would not be practical to simply start over. The current networks would need to be transitioned to support this methodology. If the network is not transitioned, we will gain a better understanding of our new functions but have a large area of functionality and traffic that is not understood

and could cause unforeseen problems. The transition needs to be considered as part of this methodology.

The development did not capture this data originally as part of the DODAF or architecture documents until the release of the DODAF v1.5 and its net-centric guidance. The data provided is a start but may not capture all the required parameters to help a network technician deploy the function to the network or maintain the function as the network evolves.

10. Conclusion

None of the concerns are insurmountable but outline the need for further research into using parameters to define the network requirements for a function. The need for improving how we incorporate new functions into the network is easy to establish. The demands are rapidly growing for consolidating networks together to serve multiple users. The traditional method of building a function, testing it on a closed network, and then deploying it to the network is no longer feasible because of the interdependencies in functions.

A more rigorous methodology needs to be implemented to reduce the impact to the operational network. The methodology needs to include a process to verify a function's network requirements before it is installed in the operational network. This can be done in a test environment or incorporated into the requirements and development.

By starting at the concept of operations, the function's network impact can be thought about as part of the functional requirements and now documented in the DODAF. This process allows the developer to build the function based on how it will be used and what it requires from the network. The tester will have something to evaluate the function against to determine if it can operate on the operational network and if not what needs to be changed to support the function.

A repeatable process will save time in deployment of the function to the operational network with minimal impact to other operational functions. The process will allow service level agreement values that are based on tested network performance. A good SLA is critical to moving forward with professionalizing the network. The service levels from the SLA will aid vendor discussions and request for proposals to identifying the right product for the requirement. Besides the SLA, the network technician will have performance numbers to give the vendor with a future network roadmap. All of this is possible because we will be using a common language to define our requirements that is shared between developers, network technicians, and the vendor community.

There are still many challenges that need to be overcome. The exact parameters need to be verified against a wide range of functions. The models need to be developed or modified to provide the necessary information needed to test the function in the test environment. Finally, the current network needs to be examined and understood to take advantage of this new methodology. Without changing the current network and deployment process, the results from using parameters will help to a certain point and then will begin to degrade because we don't have that full understanding of the network.

Although there are many challenges, the improvements will have a lasting impact to network operations. The increased capability derived from this methodology of using parameters will enhance our deployment timelines supporting SOA.

11. References

[1] DoD CIO, *Department of Defense Global Information Grid Architectural Vision version 1.0*, DoD, June 2007

[2] "Service-oriented architecture." Wikipedia. Available from en.wikipedia.org/w/index.php?title=Service-oriented_architecture&printable=yes. Internet; accessed 4 February 2008

[3] Fluke Networks, "Proven techniques and best practices for managing infrastructure changes." Fluke Network White paper. Available from <http://www.flukenetworks.com/FNet/en-us/findit?Documents=2805673>. Internet; accessed 24 January 2008.

[4] McCabe, James. *Network Analysis, Architecture, and Design*. Morgan Kaufman, 2007.

[5] Comer, Douglas. *Computer Networks and Internets with Internet Applications*. Upper Saddle River: Pearson Prentice Hall, 2004.

[6] Follows, Jonathan. "Application-Driven Networking: Class of Service in IP, Ethernet, and ATM Networks." Available from <http://www.inf.ufsc.br/~mario/QoSIBM>. Internet; accessed 14 February 2008.

[7] Menasce, Daniel, Virgilio A.F. Almeida, and Larry Dowdy. *Performance by Design: Computer Capacity Planning by Example*. Prentice Hall, 2004.

[8] Department of Defense, "DoD Architecture Framework Version 1.5 Volume I. 23 April 2007" Available from http://www.defenselink.mil/cio-nii/docs/DoDAF_Volume_I.pdf. Internet; accessed 16 March 2008.

[9] Department of Defense, "DoD Architecture Framework Version 1.5 Volume II. 23 April 2007" Available from http://www.defenselink.mil/cio-nii/docs/DoDAF_Volume_II.pdf. Internet; accessed 16 March 2008.

9. Dispensation

The views expressed in this document are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.