AD
AD-E404 131

Contractor Report ARWSE-CR-16005

COMMAND POST OF THE FUTURE BC13.2 SCALABILITY REPORT

Kiran Desilva
Michael Faulkner
William Ommert
General Dynamics C4 Systems
Aberdeen Proving Ground, MD

Ross Arnold
Michael Dykstra
U.S. Army CCDC Armaments Center
Picatinny Arsenal, NJ

September 2019

U.S. ARMY COMBAT CAPABILITIES DEVELOPMENT COMMAND ARMAMENTS CENTER

Weapons Software Engineering Center

Picatinny Arsenal, New Jersey

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-01-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (*From - To*)
September 2019	Contractor Report	

4. TITLE AND SUBTITLE
Command Post of the Future BC13.2 Scalability Report

5a. CONTRACT NUMBER
5b. GRANT NUMBER
5c. PROGRAM ELEMENT NUMBER
5d. PROJECT NUMBER
5e. TASK NUMBER
5f. WORK UNIT NUMBER

6. AUTHORS
Kiran Desilva, Michael Faulkner, and William Ommert - General Dynamics C4 Systems

Ross Arnold and Michael Dykstra - U.S. Army CCDC AC

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
U.S. Army CCDC AC, WSEC
Fire Control Systems & Technology
Directorate (FCDD-ACW-FM),
Picatinny Arsenal, NJ 07806-5000

General Dynamics C4 Systems
6245 Guardian Gtwy # 106,
Aberdeen Proving Ground,
MD, 21005

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army CCDC AC, ESIC
Knowledge & Process Management Office (FCDD-ACE-K)
Picatinny Arsenal, NJ 07806-5000

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Contractor Report ARWSE-CR-16005

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The Command Post of the Future (CPOF) has been the U.S. Army's most widely used Command Post system over the past decade. The latest version of CPOF software includes a capability that allows it to scale to over 5,000 simultaneously concurrent users. This capability allows all 5,000 users to use the CPOF software at the same time, interacting on the same data cooperatively, sharing visualizations, map-boarding, and collaborating in real time. To test this capability, a large number of computers and hardware was required. This report describes the scalability test that exercised the capability to scale live collaboration across 5,000 users. This test required a great deal of both time and monetary resources. The test was successful, proving the scalability capability of the CPOF BC13.2 software.

15. SUBJECT TERMS
CPOF Command Post CPCE CPCE v3 TacApps Command Post Computing Environment Common Operating Environment COE COE v3 Tactical Applications Scalability

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Ross D. Arnold
U	U	U	SAR	31	19b. TELEPHONE NUMBER (Include area code) (973) 724-8618

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

CONTENTS

FIGURES

FIGURES
(continued)

TABLES

ACKNOWLEDGMENTS

Weapons and Software Engineering Center (WSEC) and General Dynamics Mission Systems (GDMS) would like to thank the following organizations and individuals for their support during this scalability test effort:

- Army Test and Evaluation Command (ATEC)
- Communications-Electronics Research, Development and Engineering Center (CERDEC)
- Electronic Proving Ground (EPG)
- Program Executive Office Command Control Communications-Tactical (PEO C3T)
- Program Manager Mission Command (PM MC)
- Product Manager Tactical Mission Command (PdM TMC)
- Software Engineering Institute (SEI)
- Threat System Management Office (TSMO)
- U.S. Army Training and Doctrine Command (TRADOC)

- Jim Bell
- Brian Martinet
- Steven Miller
- Michael Phillips
- John Raletz
- Dave Smith
- Mike Wee

INTRODUCTION

At the request of Product Manager Tactical Mission Command: The Electronic Proving Ground (EPG), the Weapons and Software Engineering Center (WSEC) at the U.S. Army Combat Capabilities Development Command Armaments Center at Picatinny Arsenal New Jersey, and General Dynamics Mission Systems (GDMS) designed and executed a scalability test of the Command Post of the Future (CPOF) version BC13.2. The objective of the test was to verify that CPOF BC 13.2, based on the Third Generation (3G) CoMotion Architecture, meets the Maneuver Control System Capability Production Document requirement for a single instance of a collaborative Command and Control (C2) system to support 5,000 simultaneous users. The CPOF BC 13.2 is the first version of the system to be based on the 3G Architecture and was used in this Scalability Test. The 3G Architecture is a key part of the CPOF 13.2 software that will be fielded for Common Operating Environment, Version 1.

Two aspects of the CPOF BC 13.2 system were evaluated as part of the Scalability Test:

- System Stability – The system must remain stable when being used operationally. It must not crash under normal load and should be robust to periods of exceptional load. System stability is measured by looking for any parts of the system that fail during the test.

- System Responsiveness – The system must continuously support synchronous collaboration among ad-hoc sets of simultaneous users. In operational terms, different organizations and echelons must be able to share and interact with data (e.g. units, events) and visualizations (e.g., maps) for a portion of the mission space. A brigade (BDE), for example, must monitor and interact with a battalion (BN) to provide fire support. System responsiveness is measured by the amount of time it takes for a change injected into the system to arrive at another point in the system.

The 3G Architecture Scalability Test was a success. The system remained stable under the operational workload of 5,000 simultaneous users. During the Scalability Test, over 450 hrs of test time were accumulated on the system. During the test of record alone, the CPOF servers under test ran for over 8,000 machine-hours without failure.

The Scalability Test was executed in three phases: the test of record, excursions, and Threat System Management Office (TSMO) test environment evaluation. The testing environment was updated between each phase of the test. The CPOF BC 13.2, the system under test, was not modified after the beginning of the test of record.

During the test of record, 33 of 36 test measurements met the System Responsiveness Metrics (SRM). The three measurements that exceeded the SRM thresholds during Phase 1 were from test case 12, the largest and highest intensity test case. Analysis of the test environment logs indicated that access to disk input/output (I/O), an issue unique to the test environment, was adversely affecting all test results. The test environment was reconfigured between Phases 1 and 2 and between Phases 2 and 3 in an effort to address the disk I/O issue. Phase 3 (TSMO test environment evaluation) execution of test case 12 easily met the system responsiveness thresholds.

The Scalability Test clearly demonstrated that the CPOF BC 13.2 system, based on 3G, meets the U.S. Army TRADOC-provided requirement that a single instance of a collaborative C2 system supports 5,000 simultaneous users.

SYSTEM OVERVIEW

System Design

The CPOF is a C2 visualization and synchronous collaboration system. A synchronous collaboration system allows users located in distributed command posts to process live data, visualize diverse information, and collaborate on operations in near real time, substantially reducing the lag inherent in asynchronous collaboration systems such as email and static data systems.

The CPOF originated in a Defense Advanced Research Projects Agency program focused on advanced user interface design for C2 environments. CPOF is built on the CoMotion platform, which was derived from visualization research on the System for Automated Graphics and Explanation (ref. 1) and Visage (ref. 2).

Three design concepts lie at the heart of CoMotion: data liveness, direct manipulation, and deep collaboration.

- Data Liveness - In any C2 environment, the ability to incorporate new information dynamically is critical to the success of an operation. Though many visual analytics tools operate on static data dumps, CPOF's "live" visualizations continually update in response to changes sourced from user interactions or from underlying data feeds. The CPOF is highly composable, permitting users to author new information, integrate updated information in visualizations, or to create composite work products by assembling multiple visualizations in a single "product" container.

- Direct Manipulation – The CPOF makes heavy use of direct manipulation gestures (e.g., drag-and-drop) to afford users content management, editing, and the ability to view control operations. Simplicity and predictability emerge by employing a small set of interactions with great consistency.

- Deep Collaboration – The CPOF offers a deep collaboration capability beyond pixel sharing and chat. Any visualization or composite product in CPOF allows simultaneous interaction by every user with access to it by supporting the collaborative creation of plans and analysis products. Shared visibility among distributed team members occurs as a natural side effect of user activities.

The CPOF is used daily in distributed command posts and forward operational bases. The software spans organizational echelons from corps to battalion with users in functional areas that include intelligence, operations planning, civil affairs, engineering, and ground and aviation units. The CPOF is used extensively to support C2 operations for tasks covering information collection and vetting, situation understanding, daily briefings, mission planning, and retrospective analysis. A detailed description of CPOF's operational utility is provided in reference 3.

System Components

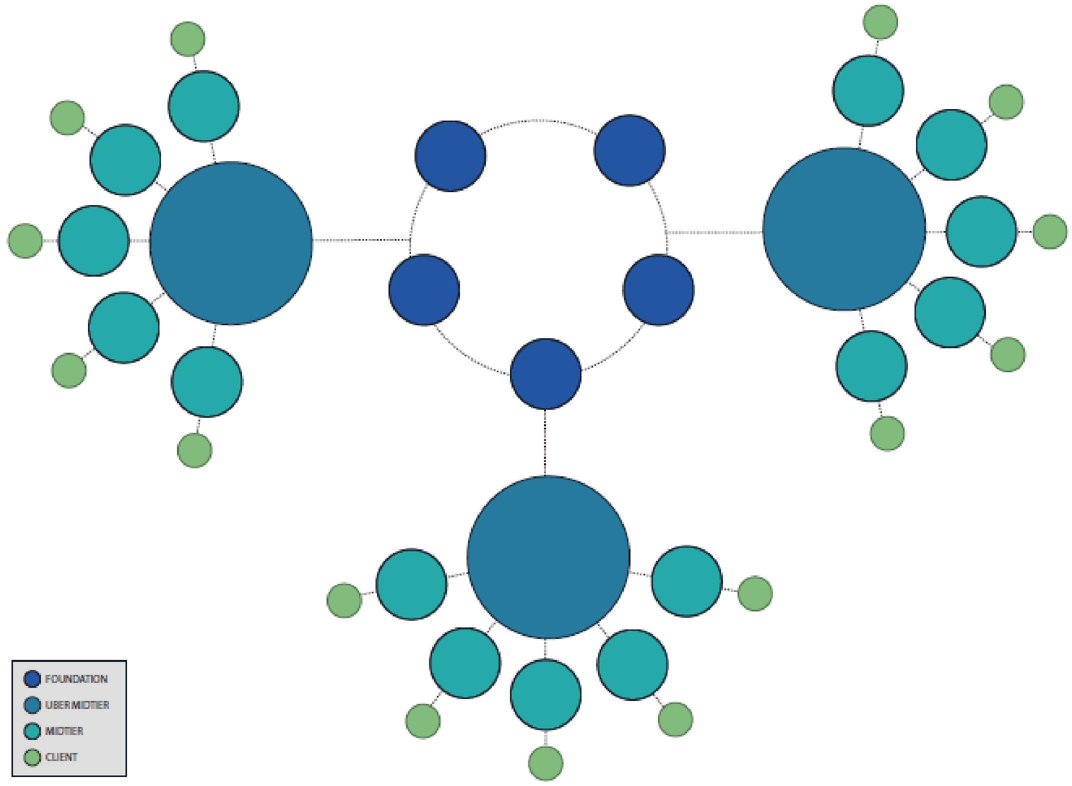

Figure 1
Sample 3G deployment architecture

The CPOF system is made up of three components:

- Foundation Servers – The foundation ring, made up of a variable number of foundation servers, hosts the primary copy of the 3G repository data, which contains the total set of information being managed by the CPOF system. Users of the same repository are operating on a common set of data. The Foundation Servers are considered to be at the “center” of the system. Traffic moving toward the Foundation Ring is moving “inward” while data moving away from the Foundation Ring is moving “outward.”

- Midtiers and Uber-midtiers – A midtier is a server that accepts connections from outward system components (Midtiers and clients) and processes changes locally. Midtiers are typically used for two reasons:

- To protect the network from unnecessary traffic. A Midtier hosted in a command post is responsible for applying and disseminating changes to all connected users. This means that a change is accessible with low latency to local users. The Midtier also ensures that each piece of data is sent over the inward connection only once. This is important since the Midtier inward connections are frequently wide area networks with significant latency and reduced bandwidth availability. If a change is made to a graphic at division and that graphic is being used by 10 users at brigade (BDE), the change will be sent to the local BDE Midtier once and the BDE Midtier will disseminate it to all connected clients.

- To provide a continuous operations capability to a set of users. All users connected to a Midtier will continue to have synchronous collaboration capabilities regardless of the connection status of the Midtier to its inward server.

The sole distinction between a Midtier and an uber-midtier is one of computing and storage resources allocated to the Midtier software process. Midtiers and Uber-midtiers are exactly the same piece of software. Uber-midtier servers are allocated substantially more computing and storage resources than a Midtier because of their more inward location and higher workload resulting from monitoring and managing more data than Midtiers deployed further outwards.A client provides the user interface that is used to plan, execute, and manage operations.

TEST OBJECTIVES AND METRICS

The purpose of the CPOF BC 13.2 Scalability Test was to demonstrate that the system is capable of supporting 5,000 simultaneous users on a single collaborative C2 system. The CPOF BC 13.2 is based on the 3G CoMotion (3G) architecture. The 3G Architecture was created to provide two primary operational benefits: continuous operations and increased operational scale (fig. 2). The continuous operations capabilities of the CPOF BC 13.2 system were evaluated in four network integration events. The increased operational scalability of the CPOF System's 3G Architecture was evaluated in this Scalability Test. While versions of CPOF based on the Second Generation Architecture have repeatedly supported 350 simultaneous users in active military operations, the 2G Architecture requires that each division's system is independent from other division or corps systems. Independently fielded systems create operational information and capability barriers along each fielded system's boundaries. The 3G Architecture allows the same information to be used by a much larger number of operational organizations without having to traverse an intermediate system or create copies of critical data.

What Motivated 3G – Why is it important?

3G enables operational goals - collaboration scale up to thousands of users and continuous operations over tactical networks

Large scale rich collaboration	Extends real time, rich collaboration to a much larger scale. Capacity for more users, more data, and more services.
Seamless individual and enclave disconnect/reconnect	Ability for users and groups to work when disconnected and to merge their changes upon reconnecting. Eliminate data management TTPs.
Collaboration across organizational boundaries	Low-echelon collaboration for cross-boundary operations (DIV1-BN2 to DIV2-BN2).
Maximum reliability	Resilient to the inevitable network and hardware failures that occur in a large tactical system. Eliminate single points of failure.
Work well over tactical networks	Preserve the efficient bandwidth characteristics of 2G while increasing flexibility. Adaptive communication for any unit, any network.
Continuous operations in changing environments and organizations	Modular BDE can seamlessly move from one DIV to another, retaining data products.
Full spectrum operations	Dynamic, fault-tolerant configurations supportive of mobile defensive and offensive operations.

Figure 2
Operational motivation for the 3G Architecture

Two aspects of the 3G Architecture were evaluated as part of the Scalability Test:

- System Stability – The 3G Architecture must remain stable when being used operationally. It must not crash under normal load and should be robust to periods of exceptional load. System stability is measured by looking for any parts of the system that fail during the test.

- System Responsiveness – The 3G Architecture must continuously support collaboration among ad-hoc sets of simultaneous users. In operational terms, different echelons must be able to synchronously share and interact with data (e.g., units, events) and visualizations (e.g., map) for a portion of the mission space. For example, a BDE must monitor and interact with a BN to provide fire support. System responsiveness is measured by the amount of time it takes for a change injected into the system to arrive at another point in the system.

Several factors influence what humans consider acceptable response times: expectations based on past experience, user tolerance for delays, and user willingness to adapt their working style to accommodate different response times (ref. 4). The optimum acceptable response time for simple tasks is 2 sec (ref. 5) and 4 sec for transaction-oriented systems (ref. 6).

The CPOF system is used for both synchronous and asynchronous collaboration. Synchronous collaboration provides the most demanding threshold for system response time. While 2 sec is considered ideal, the warfighting community using CPOF is well aware of the latencies in tactical networks. Response times of 3 sec are acceptable, and delays of up to 5 sec are tolerated. When response time exceeds 5 sec, the cost of synchronous collaboration outweighs the operational benefit.

This Scalability Test makes a distinction between holistic responsiveness and CPOF system responsiveness. Holistic Responsiveness is a function of network delays and the total processing time required by all CPOF system components (CPOF system responsiveness).

Holistic Responsiveness = Network Latency + CPOF System Responsiveness

The Scalability Test measures CPOF system Responsiveness. Characterizing the network latency component of Holistic Responsiveness requires a deployment architecture. The intended deployment architecture for a large-scale 3G-based system is tiered (fig. 1). The Foundation Ring is housed in an enterprise data center owned by a U. S. Army Service Component Command or higher organization. The Foundation Ring is not intended to support clients. Its role is to house all Repository data for a system instance. Uber-midtiers are placed with corps and Division Command Posts/Headquarters. Each Uber-midtier supports all local clients as well as any outward Midtiers. Midtiers are hosted with BDEs and possibly with BNs resources permitting. Each Midtier supports any local clients as well as any outward Midtiers or clients.

To convert the deployment network latency into network latency for the Scalability Test, a few assumptions about the 3G system network connectivity must be made:

- All communication between clients and the BDE Midtier is over a local area network (LAN).
- The 3G Midtier-to-Midtier and Midtier-to-Foundation Ring communication is over a WAN.
- All Foundation Ring communication is over a LAN.
- The LAN communication is nearly instantaneous.

- All WAN communication is satellite-based with 350 ms latency and a round-trip latency of 700 ms.

With these network assumptions, holistic responsiveness can be expanded as follows:

Holistic Responsiveness = (No. of WAN links * 350 ms) + CPOF System Responsiveness

Given the deployment architecture described previously, three specific SRMs were defined for the 3G Architecture Scalability Test:

- SRM no.1 (BDE Local): When the data is injected by a client connected to a BDE Midtier, the amount of time required for that data to arrive at a client connected to the same BDE Midtier. Figure 3 shows this.

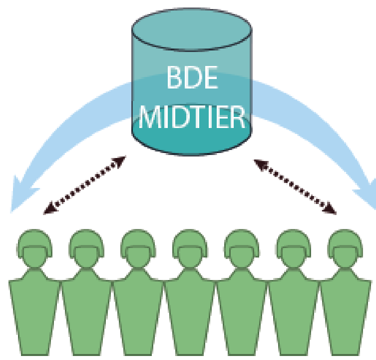

Figure 3
System traversal for SRM no.1

- SRM no. 2 (DIV Local): When data is injected by a client connected to the BDE 1 Midtier, the amount of time required for that data to transit the BDE 1 Midtier, the DIV A Uber-midtier, the remote BDE 2 MT, and arrive at the remote BDE 2 client (fig.4).

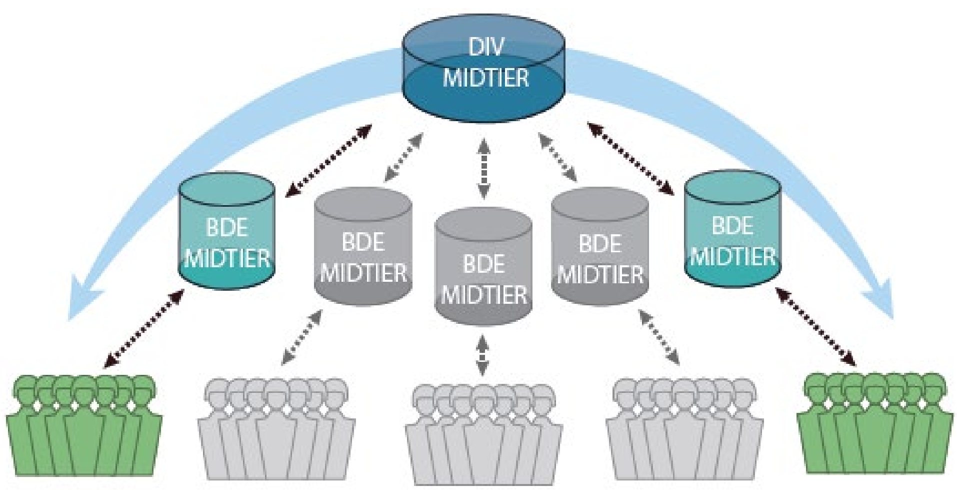

Figure 4
System traversal for SRM no.2

- SRM no. 3 (Ring Traversal): When data is injected by a client connected to the local BDE 1 Midtier, the amount of time required for that data to transit the BDE 1 Midtier, the DIV A Uber-midtier, the Foundation Ring, the DIV B Uber-midtier, the BDE 3 Midtier, and arrive at the BDE 3 client.

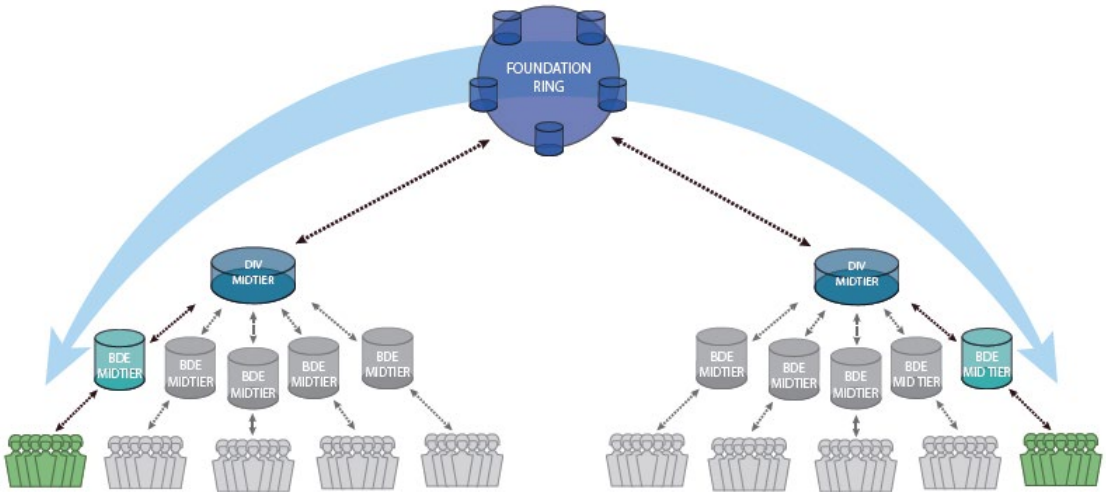

Figure 5
System traversal for SRM no. 3

In sum, based on the literature and operational experience with the CPOF system, acceptable holistic responsiveness falls in the range of 0 to 5 sec (0 to 5,000 ms). If the network travel times are subtracted from the holistic responsiveness times, acceptable CPOF system responsiveness times must fall within the ranges shown in table 1.

Table 1
Acceptable ranges for CPOF system responsiveness (in ms)

Holistic responsiveness	SRM no. 1: BDE Local		SRM no. 2: DIV Local		SRM no.3: Ring Traversal	
	Network latency	CPOF system responsiveness	Network latency	CPOF system responsiveness	Network latency	CPOF system responsiveness
2,000	0	2,000	700	1,300	1,400	600
5,000		5,000		4,300		3,600

In other words, the further the data must travel through the system, the less time the CPOF 3G system has to process and display this data. Historically, the CPOF system has been held to a stringent level of responsiveness. The 2G-based CPOF system used SRM2 (DIV Local) with a threshold of 550 ms. To be consistent with the scalability testing done on the 2G-based CPOF system, a threshold of 550 ms for SRM2 was used for this scalability test. This was then scaled down for SRM1 (BDE Local) to 250 ms and scaled up for SRM3 (Ring Traversal) to 800 ms. The first two thresholds (SRM1, SRM2) are substantially lower than the derived requirements shown in table 1. The threshold for SRM3 is in the low end of the derived range, exceeding the derived minimum by only 200 ms.

TEST MODEL

The model for the Scalability Test has two related aspects: workload and processing time. Workload is the amount of work that must be done by the system per unit of time. Processing time (referred to as "system response time") is the amount of time it takes for data to be disseminated across the 3G Architecture. This Scalability Test varied the total system workload, both the number of simultaneous users and the intensity (operations/second) of each concurrent user. System stability for each workload was monitored and System Responsiveness was measured (table 2).

Table 2
Summary of workload intensity scenarios

	Low intensity 4 ops/sec	Medium intensity 18 ops/sec	High intensity 32 ops/sec
Pilot 495 users	Test 1	Test 4	Test 7
Case 1 1518 users	Test 2	Test 5	Test 8
Case 2 2904 users	Test 3	Test 6	Test 9
Case 3 5016 users	Test 10	Test 11	Test 12

These intensity levels (operations/second) for each user are based on measurements and analysis of two operational employments of the 2G Architecture based CPOF system:

- During a BDE engagement in Fallujah (Iraq) in November 2004, users performed the equivalent of 32 3G operations per second.
- During a division engagement in the troop draw-down in Afghanistan in June 2011, users performed the equivalent of four 3G operations per second.

The Fallujah operation was the most intense use of the CPOF system to date for which data is available and is the basis for high intensity. The draw-down in Afghanistan is used as the basis for low intensity. The medium intensity is an average of High Intensity and Low Intensity.

Model Construction

Using the operational data points from Fallujah and Afghanistan as the basis for the model is valuable; however, it is important to recognize the limitations of the model. First, the model is based on a small number of measured data points. The High Intensity rate is based on a single BDE (47 users). The measured rate is linearly projected by two orders of magnitude in this Scalability Test.

It is unclear that linearly projecting from 47 to 5,000+ users provides a realistic workload. This may overstate the total system workload intensity for two reasons:

- Using a linear projection means that every system user is assigned the same activity level as each BDE user in Fallujah. In other words, for this scalability test, the 5016-user test was modeled as 19 division and corps organizations with 264 users each, producing the Fallujah engagement workload intensity.
- In a deployed environment, 5,000 users concurrently operating at the Fallujah level of intensity is unlikely. Every user is not going to operate with the same usage

profile as the core battle staff planning, executing, and managing missions in a BDE organization. For example, higher echelon organizations will have more staff assigned (Chaplain, Weather Officer, etc.). Workload intensity will vary widely depending on the user's role in the organization.

Figure 6, which is based on the numbers in table 3, shows how the workload (system operations/second) increases using a linear projection from 47 users to 5,016 users. The dotted line indicates the expected level of intensity in a deployed environment. The Expected Intensity workload adds work for each additional user, but the amount of work it adds is sublinear, taking into account the concerns expressed previously.

Table 3
Linear projection of workload per user

		Workload (operations/second)		
	No. of users	Low intensity	Medium intensity	High intensity
	1	4	18	32
	47	188	846	1,504
Pilot test	495	1,980	8,910	15,840
Case 1	1518	6,072	27,324	48,576
Case 2	2904	11,616	52,272	92,928
Case 3	5016	20,064	90,288	160,512

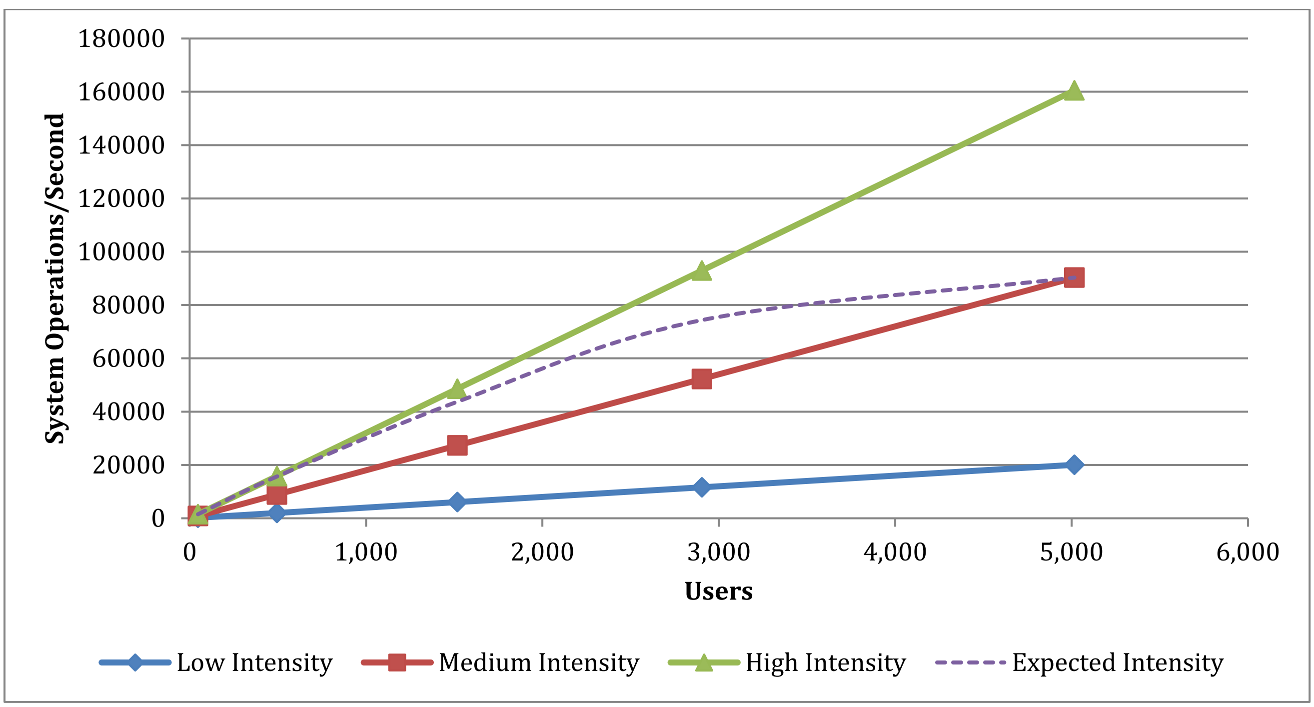

Figure 6
Linear extrapolation of system workload intensity versus expected workload intensity in a deployed environment

Second, the model uses 3G Architecture "equivalent operations" vs. the historical 2G Architecture operations to account for the fact that the 3G Architecture does more work per user action than the 2G Architecture. In the context of the Scalability Test, the most important difference is the way in which the 2G and 3G Architectures manage changes to data.

The 2G Architecture propagates all changes inward to the master server (equivalent to a foundation server in 3G Architecture) where they are managed via a centralized truth model. For example, a conflict can be generated when one user moves an enemy unit 4 km to the south at the same time another user moves the same unit 6 km to the southeast. In 2G, conflicts are resolved using a "last change wins" paradigm, where the most recent change to the data is considered the "truth copy."

The 3G Architecture requires a more sophisticated change management strategy for a number of reasons. The 3G Architecture supports continuous operations in disconnected, intermittent, and latent (DIL) environments much more robustly than the 2G Architecture. An entire BDE, for example, may lose connection to the division for a period of 2 hr. When the connection is re-established, the CPOF system must process any conflicts between the data changed by the brigade users and the users from across the rest of the system. In a distributed system running on latent networks with a large number of simultaneous users, conflicting changes to the data, such as the one described previously, are inevitable. The rate of conflicts depends on the organization's standard operating procedure and tactics techniques and procedures in using the system. The CPOF 3G Architecture processes conflicting changes using the stored change history of the data.

The 3G Architecture stores data in terms of its change history rather than just its current value. Conflicts are detected and represented explicitly, eliminating the problem of unintentionally overwriting the work of others. Representing data in terms of its change history allows the system to preserve and correctly resolve conflicts, whether a brigade is reconnecting after 2 hr or whether two users concurrently made conflicting changes to the same unit (figure 7).

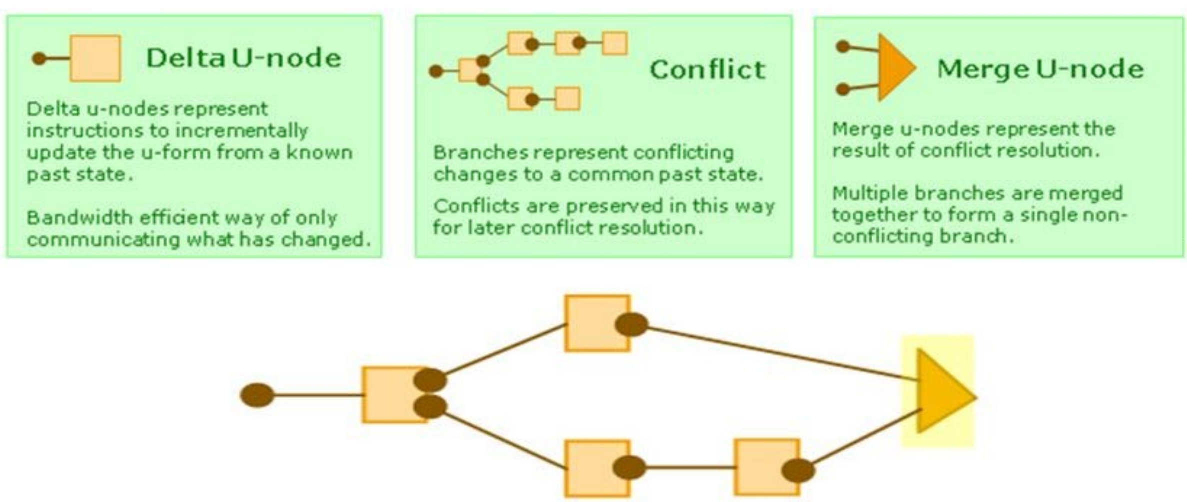

Figure 7
3G Architecture conflict resolution mechanism

Managing the change history increases the workload on the system, which affects scalability and performance. Consequently, the testing scenarios include conflict resolution in the workload.

Differences between Operational Provisioning and Test Environment

Any specific hardware's ability to support a given workload intensity is fixed. If more operations are coming in per second than what the integrated system (hardware and software) can handle, the system will eventually slow down and ultimately become unstable. The 3G Architecture, as developed and delivered by GDMS, is a pure software solution. When deployed and used, the 3G Architecture is run on computing hardware through the use of virtualization (e.g., VMWare virtual machines). As a software system, the 3G Architecture's performance is sensitive to the hardware resources provisioned to support it.

When operationally employed, 3G Architecture system components (Foundation Servers, Midtiers, and clients) are provisioned to have local computing resources (central processing unit CPU and memory) as well as local storage (disk via filers). Each Command Post BDE and above will have hardware provisioned for hosting 3G Architecture server processes (Foundation and Midtier). The server virtual machines have typically been hosted on the Battle Command Common Server (BCCS). At times, the Intel Fusion Server has hosted and selected 3G Architecture Server Processes and the Tactical Server Infrastructure is anticipated to host these processes in the future.

System under Test

The provisioning for the 3G Architecture-based CPOF BC 13.2 System Components under test is shown in table 4.

Table 4
Hardware specifications for a fielded CPOF BC13.2 system

			Disk Space	
System component	Virtual CPUs	Virtual random access memory (RAM)	OS Disk (C:)	Data disk (E:)
Foundation server	12	80 GB	60 GB	250 GB
Uber-midtier (DIV and CORPS)	12	80 GB	60 GB	250 GB
Midtier (BDE)	4	32 GB	60 GB	50 GB

The CPOF BC 13.2 clients are Java Thick clients that are typically hosted on laptops (no virtualization). The clients leverage the laptop resources (CPU, memory, disk space). Each client runs with 4 GB of RAM and 465 GB of local disk storage. CPOF client software running on Warfighter Information Network – Tactical uses operating system level virtualization to host the software on a point of presence. In this case, the hardware provisioning is reduced.

All of the server and client processes are run in 64-bit mode on 64-bit Windows Operating Systems. The core 3G Architecture components have little reliance on the Windows Operating System and are portable to other environments (e.g.,*nix including Linux and Android).

Other Scalability Test Components

Performance testing is a proxy for usage in the real world. Getting 5,000 concurrent users on a 3G system was impractical in terms of computing, network, and human resources. Therefore, for the Scalability Test, almost all of simultaneous users were simulated via load generating software. In addition to the simulated users, certain tests included the use of a small number of human-operated BC 13.2 CPOF clients operating against the system under “test as additional simultaneous users.”

In order to carry out the Scalability Test, several other components were required. The provisioning for these components can be found in table 5.

Table 5
Hardware specifications for additional scalability test components

			Disk space	
System component	Virtual CPUs	Virtual RAM	OS disk (C:)	Data disk (E:)
Load generators	6	24 GB	60 GB	0
Domain controller	4	16 GB	60 GB	0
STAF/STAX server	16	128 GB	60 GB	8000 GB
*Log aggregation	20	32 GB	60 GB	1000 GB
*EPG net forensics	4	32 GB	60 GB	1000 GB

Each load generator executed 33 3G cores in order to simulate 33 clients generating load on the servers under test. All 33 of these cores were running within a single Java virtual machine and each 3G core was configured to use an in-memory Repository to permit a reduction in disk I/O requirements by a factor of two. The domain controller was a server that allowed the clients to access the system by providing Dynamic Host Configuration Protocol, domain name server, and Domain authentication services. The Software Testing Automation Framework/Software Testing Automation Engine (STAF/STAX) server was the component that coordinated the execution of all of the scalability test cases, starting all clients and servers, directing them to perform work, and gathering and analyzing the test results. The log aggregation server stored the test logs generated by the 3G system servers and clients including test results in tab separated value files (.tsv) machine performance statistics and 3G system logs. The Electronic Proving Ground Net Forensics server monitored the network traffic between servers during all test ru.

Difference between Operational and Test Hardware

While the test environment was an excellent proxy and contained resources to run the test, it differs from the expected deployment environment in two significant ways:

- Access to the Disk I/O:In a deployed environment, one Storage Area Network (SAN) device would be available to every two blades in a BCCS server stack. The test environment was set up to simulate a deployment of 19 divisions. Each division is provisioned to have 10 BCCS server stacks, or 20 blades, with access to 10 SANs (NetApp filers). In a deployed environment, a 5,000-user CPOF BC 13.2 system would have 190 NetApp filers. In the test environment, all storage was centralized on only two large NetApp filers.

- Network Characteristics: The CPOF BC 13.2 system is designed to operate reliably in (DIL) tactical networking environments that are notoriously latent and prone to high bit-error rates. The test environment ran using Enterprise networking assumptions – low latency and low bit-error rates. A tactical network would naturally buffer the server processes by delivering changes at a lower rate than an enterprise network. The scalability test can also be viewed as a form of 3G Architecture stress test.

Test Overview

Table 6 shows the number of user-based clients, load generators (simulated clients), the total number of clients under test, and the number of BDE, DIV, and foundation servers under test.

Table 6
Total Number of Servers and Clients per Test Case

Test Case	Physical Client Count	Load Generators (33 virtual clients each)	Active Clients	Brigade Midtier (MTier)	Division Midtier (UTier)	Foundation Server Count
Test 1 (Pilot, Low Intensity)		15	495	15	3	3
Test 2 (Pilot, Medium Intensity		15	495	15	3	3
Test 3 (Pilot, High Intensity)	4	15	499	15	3	3
Test 4 (Case 1, Low Intensity		46	1518	46	6	9
Test 5 (Case 1, Medium Intensity)		46	1518	46	6	9
Test 6 (Case 1, High Intensity	4	46	1522	46	6	9
Test 7 (Case 2, Low Intensity)		88	2904	88	11	17
Test 8 (Case 2, Medium Intensity		88	2904	88	11	17
Test 9 (Case 2, High Intensity)	4	88	2908	88	11	17
Test 10 (Case 3, Low Intensity)	4	152	5020	152	19	34
Test 11 (Case 3, Medium Intensity)	4	152	5020	152	19	34
Test 12 (Case 3, High Intensity)	4	152	5020	152	19	34

Some tests used only virtual clients (33 per load generator). Other tests used the virtual clients plus active thick clients operated by human users. During the test of record phase of the Scalability Test execution, for example, multiple physical CPOF clients operated by users were connected to the Midtiers. An operator created a map and shared it with another operator to establish a collaborative session. Once the collaborative session was established, a unit was created and dropped on a collaborative map. The time taken for the other operator to receive the unit on the other map was measured. In addition, the color and location of the unit was updated, and the amount of time taken to receive the change by a collaborating user was measured.

TEST OF RECORD RESULTS

Multiple runs of each test case were executed. The system response time for each test was measured between clients in the same BDE, the same division, or the same theatre at each level of workload intensity. Each test case was measured against these SRM.

- SRM No.1 (BDE Local): ≤ 250 ms
- SRM No.2 (DIV Local): ≤ 550 ms
- SRM No.3 (Ring Traversal): ≤ 800 ms

The results in table 7 indicate elapsed time in milliseconds from when an update was submitted at one client venue to when that change was observed at a second client venue. These numbers represent the set of measurements with the best ring travel for each test case.

Table 7
Test of record Results (green = pass, yellow = fail)

	SRM1: BDE local (ms)	SRM2: DIV local (ms)	SRM3: ring traversal (ms)
Test 1 (Pilot, Low Intensity)	4 ms	7 ms	74 ms
Test 2 (Pilot, Medium Intensity	5 ms	9 ms	14 ms
Test 3 (Pilot, High Intensity)	4 ms	10 ms	164 ms
Test 4 (Case 1, Low Intensity)	5 ms	8 ms	14 ms
Test 5 (Case 1, Medium Intensity)	4ms	11 ms	18 ms
Test 6 (Case 1, High Intensity)	5 ms	21 ms	42 ms
Test 7 (Case 2, Low Intensity)	5 ms	9 ms	15 ms
Test 8 (Case 2, Medium Intensity	7 ms	24 ms	47 ms
Test 9 (Case 2, High Intensity)	17 ms	109 ms	236 ms
Test 10 (Case 3, Low Intensity)	5 ms	10 ms	53 ms
Test 11 (Case 3, Medium Intensity)	28 ms	202 ms	515 ms
Test 12 (Case 3, High Intensity)	614 ms	19,761 ms	19,267 ms

The 3G Architecture Scalability test of record was a success. Thirty-three of 36 test measurements were under the test threshold. Of the three test measurements, all of them observed during test 12 (Case 3, High Intensity) were over the test thresholds established in the SRMs. However, these test results were an artifact of network configuration issues discovered during testing. When those issues were corrected, all of the Scalability Test results were well within the System Responsiveness Metrics.

The Scalability Test was executed in three phases:

1. The government-witnessed test of record – this phase was designed and executed by General Dynamics with the support of WSEC, EPG, and the Software Engineering Institute.

2. Excursions – this phase was executed after the test of record and included several tests that WSEC and the TCM requested to aid in deployment decisions.

3. TSMO Test Environment Evaluation – this phase was executed by TSMO staff who used the test 12 (case 3, high intensity) to validate network configuration changes.

In preparation for the test of record, the TSMO worked to correct a problem in their network and hardware configuration. Although the problem remained, the decision was made to go ahead with the Scalability test of record. The CPOF BC13.2 software, including the 3G Architecture, went through all the test cases with no changes. The results of the test of record are shown in table 7.

During the test of record, TSMO personnel were monitoring their equipment. The SAN management software wrote log messages indicating substantial I/O latency from the SAN, and the virtual machines running the system under test reported frequent SAN disconnects. The two SANs in the test environment were configured with 3 TB of solid state drives to provide low latency read. The I/O latency issues were not with the SANs but the SANs network connectivity. The excessive I/O latency and the frequent disconnects appeared in the test of record logs as delays in reads and writes between the servers and the SAN in the I/O Analysis section. Each I/O delay pauses most server processing, leading to long processing delays for that server. Every server was impacted by I/O delays, and the occurrences of delay were random. As updates must pass through numerous servers when transiting the system, these processing delays accumulate, increasing the system response times or the time it takes for changes to propagate through the system. The SRM3, Ring Traversal, was most significantly impacted as data had to traverse at least five servers.

After the test of record, and before the excursions phase, TSMO personnel took the opportunity to upgrade the iPhone operating system on one of the Nexus 6k switch from 7.0.3 to 7.0.6. During the Excursions phase, Test 12 (case 3, high intensity) was re-run in order to gather data that would aid with deployment decisions. The read times between the servers and the SAN improved from a high of 20 sec to a high of 2.5 sec. Similarly, the times also improved from hundreds of 20-60 sec delays to fewer delays at only 10-40 sec.

After both the test of record and the excursions were completed, TSMO continued to troubleshoot the SAN I/O latency and frequent disconnects from the host machines. Each SAN had one 10 gigabit network interface for the test of record and the excursions. The TSMO increased that number to four 10 Gigabit interfaces for each SAN and re-ran test 12 (case 3, high intensity) to verify their configuration changes were a success. The read times dropped to an average of 0.02 sec and write times dropped to an average of <2 sec

To further clarify the differences in the results during each phase, the following timeline indicates the history of test 12 (case 3, high intensity) test runs. It is important to note that the software, hardware, and configuration of the CPOF BC 13.2 system under test were held constant through all testing phases. The network configuration was modified between Phases 1 and 2 and between Phases 2 and 3 as indicated previously. The final test (7/15) was executed by TSMO staff to validate their network configuration changes, and TSMO provided the results to WSEC after that phase was complete.

- Test of record Phase - June 16th, the case 3 high intensity run reported an average ring travel time of 19.27 sec.
- Excursions Phase - June 26th, the case 3 high intensity run reported an average ring travel time of 4.43 sec.
- TSMO Phase – July 15th the Case 3, High Intensity run reported an average ring travel time of 0.277 sec (277 ms), a decrease of several orders of magnitude from the test of record system response time.

This difference in travel times between the test of record Phase and the TSMO Phase would be most profound under high workloads – those with high user counts and high intensity. Therefore, if the test of record were to be repeated using the TSMO Phase network configuration, it is expected that the system response times for all 36 Test Cases would be lower with the case 2 high and case 3 medium intensity tests showing the most significant decreases in travel times. In effect, after the network configuration issues were resolved, the tests that exceeded the threshold during the test of record passed with system response times well under the SRMs (table 8).

Table 8
Test 12 case 3, high intensity for each scalability test phase

Test 12 - case 3, high intensity (times in milliseconds)	SRM1:BDE Local	SRM2: DIV Local	SRM3: Ring Traversal
test of record	614	19,761	19,267
Excursions	298	2,246	4,433
TSMO Environment Evaluation	19	128	277

ANALYSIS AND ADDITIONAL TEST RESULTS

Input/Output Performance Analysis

During test execution, one set of data that was captured was average read and write times to the SAN. The Windows Performance Monitoring utility logged the average read/write times at 1-sec intervals. The figures in this section show the average time in seconds that Foundation Servers saw for disk read and write requests to complete. Most of the time, the average times I/O access times were in the 2 to 10 ms range. However, the following figures illustrate the high disk I/O load that exposed network configuration issues during the Scalability Test.

Each of the 205 servers used during this test saw slightly different access times to the SAN based on its activity and the activity of the other servers. Therefore, the graphs included here are representative examples of the servers used in this test. While the I/O latency issues in the test environment affected all servers in the system, this analysis focuses on the foundation servers. The analysis is not specific to the foundation servers, and the findings can be applied to all servers in the system under test.

Figures 8 and 9 (June 16th log data) and, to a slightly lesser extent, figures 10 and 11 (June 26th log data) show 3 to 5 instances of lengthy read delays while accessing the SAN. Normal access time to the SAN is 2 to 15 ms, and access times over 100 ms are severe delays. As can be seen in the figures, the servers were experiencing multiple delays of as much as 20 sec on June 16th and over 2 sec on June 26th. These read delays to the SAN pause most 3G Server processing. As updates must pass through five servers when traveling through the ring, these processing delays accumulate, leading to long periods of delay until each server recovers, affecting overall system response time.

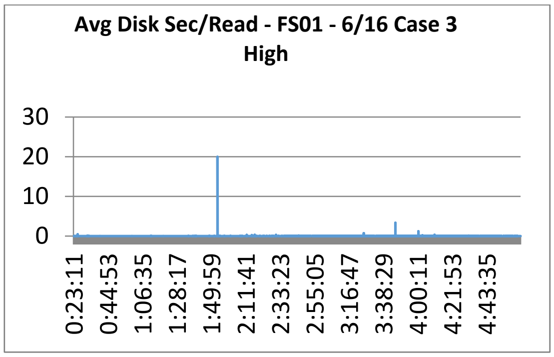

Figure 8
Read delays, foundation server 01, 6/16

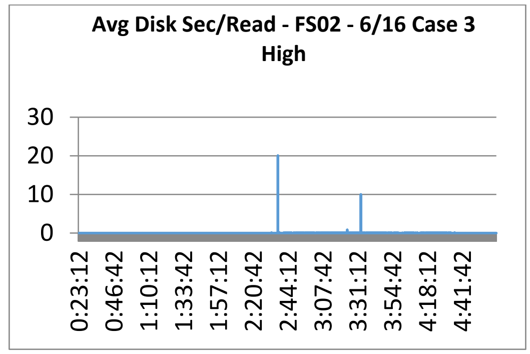

Figure 9
Read delays, foundation server 02, 6/16

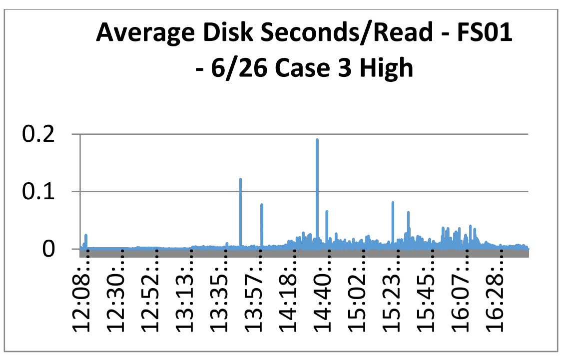

Figure 10
Read delays, foundation server 01, 6/26

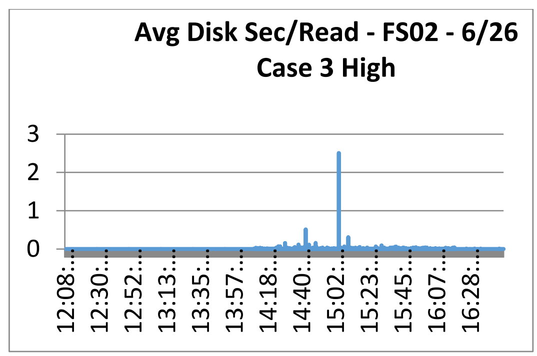

Figure 11
Read Delays, foundation server 02, 6/26

Figures 12 (June 16 log data) and 13 (June 26 log data) illustrate the much more severe and pervasive write delays. Normal processing time for a write to the SAN is 2 to 10 ms. The servers were experiencing hundreds of multi-second delays, many in the 20 to 60 sec range. The write delays were longer and more frequent; however, since the server's queue writes for asynchronous processing, there is no immediate effect on system responsiveness. It is only when the number of queued writes becomes excessive that the 3G servers slow change processing. As a result, these write delays had less immediate effect than the read delays. However, as the test progresses, these write queues get larger and are more likely to cause the servers to slow change processing in order to maintain system stability.

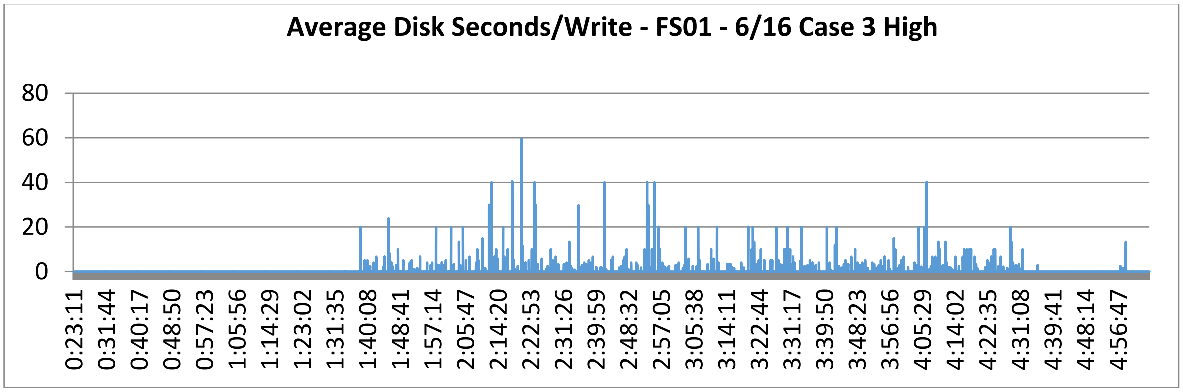

Figure 12
Write delays, foundation server 01, 6/16

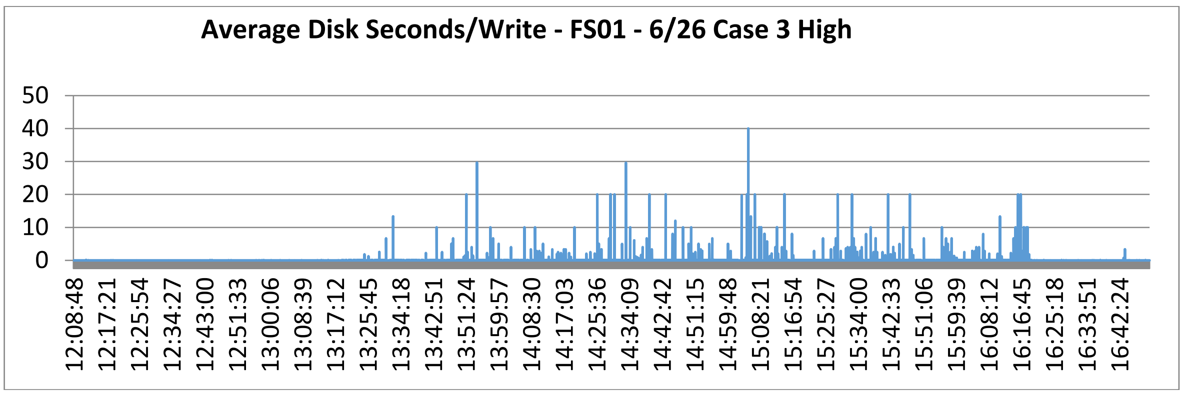

Figure 13
Write delays, foundation server 01, 6/26

As shown in the figures 12 and 13, these effects were worst during the test of record (6/16 log data) and modestly improved during the excursions (6/26 log data). There is a profound improvement seen in the TSMO testing (7/15 log data). As illustrated in figures 14 and 15, adding more network interfaces to the SAN eliminated the long read times and reduced the long write times from hundreds of reports to a handful. The corrections to the network platform made after the Excursions Phase and before the TSMO Phase allowed the data to flow smoothly to and from the SAN, resulting in profoundly improved system response times.

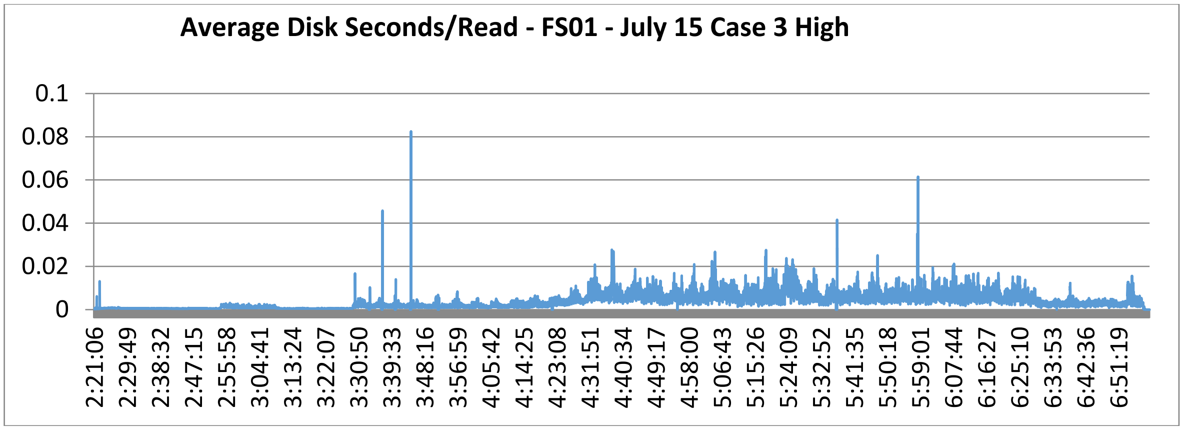

Figure 14
Read delays, foundation server 01, 7/15

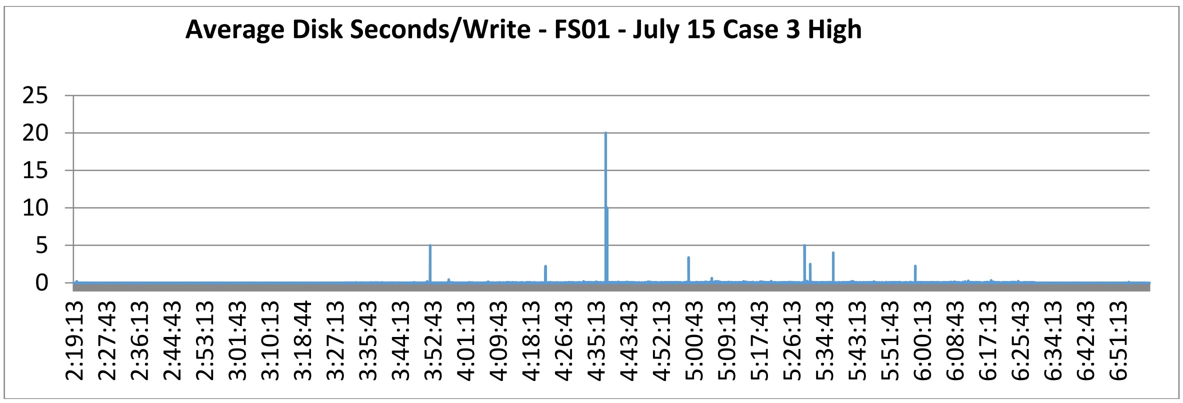

Figure 15
Read delays, foundation server 01, 7/15

In summary, the system response times for case 3, high intensity, during the test of record and the excursions is an artifact of the test environment. In the deployed environment, one SAN would be available to every two blades in a BCCS server box. In the test environment, all storage was centralized on only two large NetApp filers. If the test environment had exactly mirrored the deployed environment, the I/O issues caused by the network configuration would not have occurred, and there would not have been a negative impact on system response times. Ideally, all tests would be re-run on the corrected platform to better characterize the system performance since it is likely that the test measurements are all negatively impacted by the I/O access issues. It is expected that the system response times for all 36 test cases would show improvement with the case 2 high and case 3 medium intensity tests showing the most significant decreases in travel times.

Scalability by Echelon Analysis

The scalability test data highlights an area where an enhancement to the 3G Architecture will make the system even more robust and scalable. Part of the data captured by the 3G Architecture is venue wide distress (VWD). The 3G uses VWD to flexibly manage all the internal processes of the architecture. The VWD has values between 0 and 10 and is a measurement of the current workload of each 3G server. The 3G servers that are keeping up with all incoming changes will have a VWD of zero. The 3G servers with a VWD of 1 are experiencing workload beyond what they can process as it comes in. The 3G servers that are substantially overtaxed will have a VWD of 10. A 3G system will experience increases and decreases in VWD as the workload changes over time. A VWD of 10 is not a problem unless it is sustained. As VWD rises, the 3G core pauses noncritical functions (search indexing, repository garbage collect, and other custodial activities) to prioritize the timely processing, dissemination, and delivery of changes. A VWD of 3 or higher that is sustained for a period of 24 hrs is an indication that the venues are under-provisioned. Venues that are chronically under-provisioned will eventually become unstable as custodial services are critical to long-term system stability.

As discussed in the test model, the deployed system configuration is expected to reflect the Task Organization (TaskOrg) with the Foundation Ring housed in an enterprise environment, divisions hosting Uber-midtier servers, and BDE hosting Midtier servers. The total workload is highest at the Foundation Ring since all changes are eventually processed and stored there. The next highest workload is at the division Uber-midtiers because they process changes for all users at that echelon and also aggregate and process all changes to and from the outward BDEs. The BDEs-Midtiers have the lowest workload because they are aggregating changes for the local users only.

Currently, the foundation ring is the point at which the 3G system is most easily scaled up or down in response to number of users and total workload. If the foundation servers, for example, are consistently running at too high of a VWD, they can be added to the ring, lowering the workload required of each individual server. Figures 16 and 17 show VWD data from the case 3 high intensity test. These graphs show that the foundation servers are typically running at a low VWD but have numerous spikes up to 10. The foundation servers easily handled the workload intensity for a 3G system of the size modeled in this scalability test, and it recovered from the VWD spikes as designed.

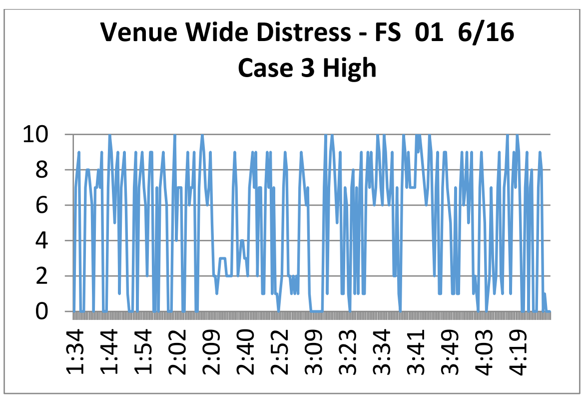

Figure 16
Foundation server distress levels, 6/16

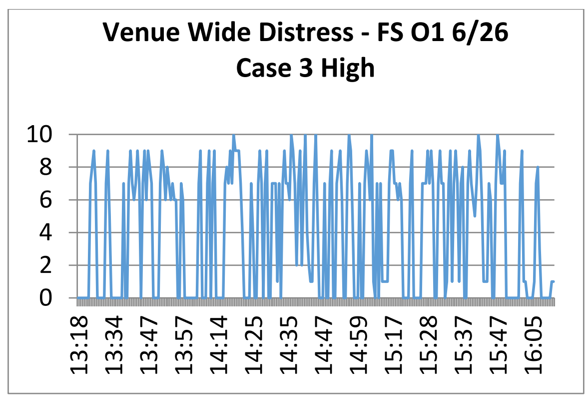

Figure 17
Foundation server distress levels, 6/26

Midtiers and Uber-midtiers can currently be scaled to meet a greater number of users and a higher workload in two ways: by provisioning existing Midtiers to handle the workload or by adding independent Midtiers. As shown in figures 18 and 19, the BDE Midtiers are provisioned with an adequate number of cores, amount of RAM, and disk space for the expected workload. The VWD is typically very low and shows occasional spikes.

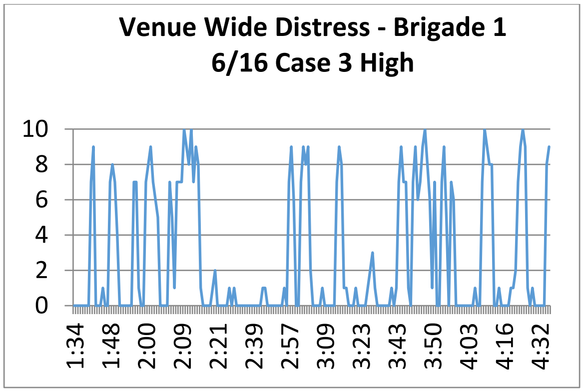

Figure 18
Brigade midtier distress levels, 6/16

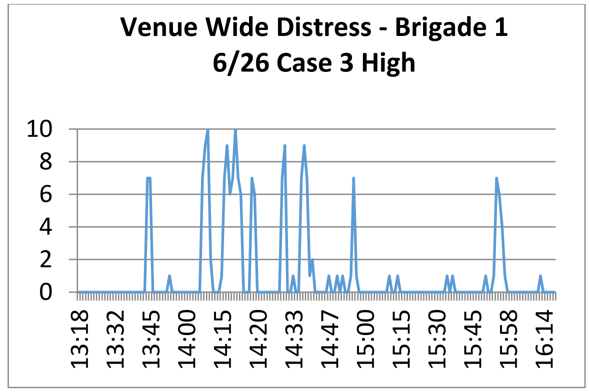

Figure 19
Brigade midtier distress levels, 6/26

Figures 20 and 21 show that the VWD for the division Uber-midtier used in this scalability test is consistently at 7 with occasional spikes above that level. While the system is working as designed, this is an indicator that the division Uber-midtiers are under-provisioned for the workload. At VWD 7, the Uber-midtier servers responded by pausing search indexing, pausing custodial activities, and slowing the rate at which they accepted new work, which slowed system response times but maintained system stability. Running the division Uber-midtier at a sustained VDW of 7 will eventually lead to system instability.

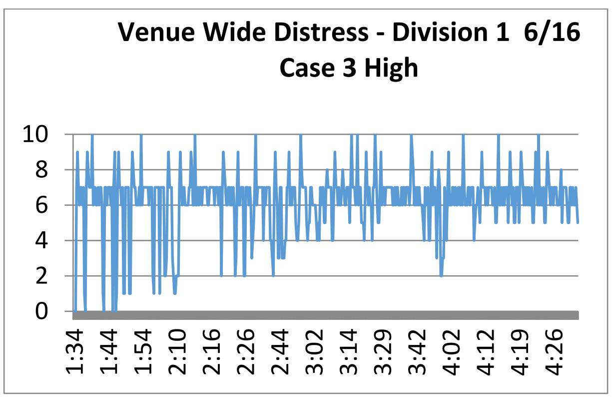

Figure 20
Brigade uber-midtier distress levels, 6/16

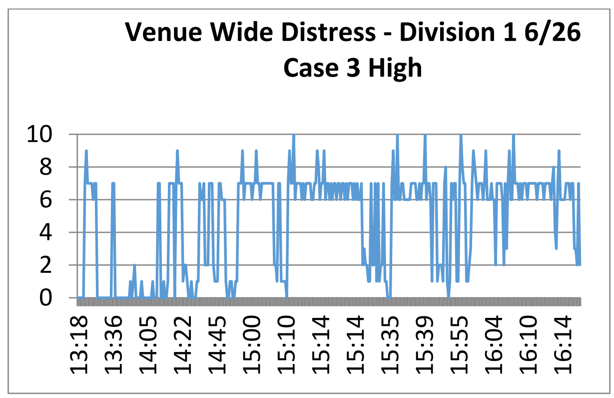

Figure 21
Brigade uber-midtier distress levels, 6/26

It is important to recall that Midtiers and Uber-midtiers are the same software component: the only difference between them is that Uber-midtiers are provisioned to have substantially more computing resources. Increasing the provisioning of the division Uber-midtier is unlikely as it is already using substantial resources. Adding an independent Uber-midtier at division can be done but has negative consequences from a system robustness perspective and may lead to increased WAN network traffic by requiring intra-division traffic to traverse to and from the ring. An enhancement that could be made to the 3G Architecture is to allow Midtiers to ring in the same way as the Foundation Servers do currently.

The 3G Architecture was designed to allow ring formation at all levels of the system, but the implementation and testing of ringing has focused only on the Foundation Servers to date. Adding the ability to ring Midtiers would allow a division to more flexibly use the computing resources allocated to it. Additionally, a ring of Midtiers could replace the Singleton Uber-midtier at Division and would increase system reliability by eliminating the single point of failure in the CPOF system at Division.

Network Conditions Analysis

The system remained robust and available over the entire course of the scalability test during even the highest workloads. As previously indicated, this scalability test was conducted over a low-latency, enterprise-class network. In a fielded environment over a tactical network, the CPOF BC13.2 3G Architecture software will operate on the tactical network. The tactical network will constrain bandwidth, which will buffer each 3G server from spikes in workload. Conducting the Scalability Test over a clean enterprise network is a type of a stress test for the system since the network did not buffer any of the 3G servers from workload spikes. The system responded as it was designed to do. As the workload intensity increased, outward pressure on the system slowed it down so that the workload intensity could be accommodated. The system did not run out of memory, crash, refuse connections, or become unavailable at any point during the test. In short, the 3G system operated reliably under the heaviest workloads possible.

CONCLUSIONS

The Third Generation (3G) Architecture scalability test was a success. During the test of record alone, the Command Post of the Future (CPOF) 3G server processes ran for more than 8,000 hours without any fault or failure.

During the government-witnessed test of record, 11 of 12 test cases (33 of 36 test measurements) were within the established system responsiveness threshold. Once issues unique to the test environment were addressed by Test System Management Office (TSMO), Test 12 (Case 3, high intensity) also met the system performance metrics. The analysis of the logs lead us to believe that the measurements for all tests are overstatements of the system responsiveness. The system responsiveness is expected to would improve for all tests if re-run in the updated test environment.

The I/O artifacts in the test environment caused the 3G Architecture to respond as designed to systemic distress. It remained stable by slowing change processing, and seamlessly recovered when the I/O blockages cleared.

The analysis of venue wide distress from the scalability test leads to recommend implementing Midtier Ring formation for 3G. When future Tactical Mission Command applications are fielded and the number of users increase, Midtier ring formation will be an important tool to increase system scale and robustness.

One of the tests requested by U.S. Army Training and Doctrine Command (TRADOC) capabilities manager for the excursions phase was to run the CPOF system over a period of 150 hr. Re-running the test over a longer period of time could further validate that the CPOF BC13.2 System 3G Architecture can robustly support 5,000 simultaneous users.

REFERENCES

1. Roth, S.F., Goldstein, J., Kolojejchick, J., and Mattis, J., "A Framework for Knowledge-Based, Interactive Data Exploration," *Journal of Visual Languages and Computing*, 5: 339-363, 1994.

2. Roth, S.F., Lucas, P., and Gomberg, C.C., "Visage: Dynamic Information Exploration," Conference *Companion of the Conference on Human Factors in Computing Systems,* 1996.

3. Croser, C., "Commanding the Future: Command and Control in a Networked Environment," *Defense and Security Analysis* 22(2):197-202, 2006.

4. Shneiderman, B., "Response Time and Display Rate in Human Performance with Computers," Computing *Surveys*, 16(3): 265-285, 1984.

5. Miller, R.B., "Response time in man-computer conversational transactions," *Proceedings of the Spring Joint Computer Conference*, 33, AFIPS Press, Reston, VA, 267-277, 1963.

6. Williams, C.M., "System response time: A study of users' tolerance," *IBM Advanced Systems Development Division Tech. Rep*. July, IBM, Yorktown Heights, NY, 17272. 1973.

DISTRIBUTION LIST

U.S. Army CCDC Armaments Center
ATTN: FCDD-ACE-K
FCDD-ACW-HH, M. Dykstra
FCDD-ACW-FM, R. Arnold
Picatinny Arsenal, NJ 07806-5000

Defense Technical Information Center (DTIC)
ATTN: Accessions Division
8725 John J. Kingman Road, Ste 0944
Fort Belvoir, VA 22060-6218

General Dynamics C4 Systems
ATTN: K. Desilva
M Faulkner
W. Ommer
6245 Guardian Gtwy # 106
Aberdeen Proving Ground, MD 21005

REVIEW AND APPROVAL OF ARDEC REPORTS

THIS IS A:

☑ **TECHNICAL REPORT** ☐ **SPECIAL REPORT**

☐ **MEMORANDUM REPORT** ☐ **ARMAMENT GRADUATE SCHOOL REPORT**

FUNDING SOURCE PM funded Program of Record
[e.g., TEX3; 6.1 (ILIR, FTAS); 6.2; 6.3; PM funded EMD; PM funded Production/ESIP; Other (please identify)]

Command Post of the Future BC 13.2 Scalability Report	TacApps / CPOF
Title	Project
Ross Arnold, Mike Dykstra	ARWSE-CR-16005
Author/Project Engineer	Report number/Date received (to be completed by LCSD)
x8618	WSF-M
Extension Building	Author's Office Symbol

PART 1. **Must be signed before the report can be edited.**

a. The draft copy of this report has been reviewed for technical accuracy and is approved for editing.

b. Use Distribution Statement A ✓, B____, C____, D____, E____, or F____ for the reason checked on the continuation of this form. Reason: public release

1. If Statement A is selected, the report will be released to the National Technical Information Service (NTIS) for sale to the general public. Only unclassified reports whose distribution is not limited or controlled in any way are released to NTIS.

2. If Statement B, C, D, E, or F is selected, the report will be released to the Defense Technical Information Center (DTIC) which will limit distribution according to the conditions indicated in the statement.

c. The distribution list for this report has been reviewed for accuracy and completeness.

Tri Lu 6/28/2017
Division Chief (Date)

PART 2. To be signed either when draft report is submitted or after review of reproduction copy.

This report is approved for publication.

Tri Lu 6/28/2017
Division Chief (Date)

RDAR-CIS (Date)

LCSD 49 (1 Sept 16)
Supersedes SMCAR Form 49, 20 Dec 06