RDECOM CERDEC US ARMY—RDECOM

Optimization of Available Bandwidth through Mission-Based Service Management

Mission Engaged Research in Cloud Use for the Army

Project MERCURY

White Paper

Communications-Electronics Research, Development, and Engineering Center (CERDEC) Command, Power, and Integration Directorate (CP&ID)

30 November 2012

DRAFT

UNCLASSIFIED

TABLE OF CONTENTS

LIST OF FIGURES

1 EXECUTIVE SUMMARY

Information superiority is the key to winning today's battles. This superiority requires timely and targeted flow of data to the right place at the right time. To make this possible, the U.S. Army is deploying more powerful and more intelligent capabilities closer to the battle front line. The movement of the data requires a robust network. The Army's network capacity and stability are improving, but bandwidth continues to be a constrained, variable and contested resource, even as new data sources are continuously added. This paper describes the project titled "Mission Engaged Research in Cloud Use for the Army (MERCURY)" – the latest approach developed by CERDEC CP&ID for dynamically allocating and optimally utilizing the available bandwidth as well as other shared Information Technology (IT) resources such as memory and processing power in the most mission-focused way.

Project MERCURY is a software capability that analyzes and manages individual web services. When these services are grouped to support a mission, a mission profile is created. A mission profile is a dynamic grouping of services needed for a single mission. Project MERCURY also plans to provide a commander-friendly dashboard for rapid balancing of the shared IT resources such as bandwidth among active mission profiles to ensure constant availability of mission-essential information. It accomplishes this by relying on a set of non-intrusive web service collection agents, which measure the resource utilization; and allocation agents, which enforce the desired allocations by means such as throttling, time shifting or intelligent routing. Initially, Project MERCURY will demonstrate the dynamic management of web services as a reaction to the changing bandwidth availability and mission focus. Future phases will address dynamic, profile-based management of other IT resources such as processing and memory.

2 INTRODUCTION

The Department of the Army has made IT and Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) system modernization and consolidation a high priority. Convergence of heterogeneous systems onto a common platform is a core element of the modernization as described by the Army CIO in "Army Software Transformation Strategy" (Army CIO/G6, 2010). Sharing of resources reduces the IT footprint but requires intelligent utilization of available IT resources. For instance, the backbone of IT and C4ISR systems is the network characterized by its bandwidth. Limited or variable bandwidth is one of the greatest information-affecting challenges in tactical environments. Allocating bandwidth effectively requires a great deal of advanced expertise, often not available in every unit, and software systems that can adjust to the changing conditions. As most units don't have either, they simply unplug systems that consume too much bandwidth, thus rendering some parts of their IT infrastructure temporarily unavailable. Project MERCURY, described in this paper, proposes a new approach to allocating limited, shared IT resources, initially focusing on dynamically managing web services to better utilize the available bandwidth. This approach will help commanders understand, budget and optimally allocate these critical resources without sacrificing essential mission capabilities or requiring specialized expertise.

The need to better utilize available bandwidth is similar to a scenario that played out in Iraq from 2003 – 2007 when the U.S. military established temporary bases and set up basic electrical services for the soldiers. Initially, the power distribution was limited to supporting internal lighting and notebook computers; however, as time passed, soldiers found the power supply stable, so they added coffee makers and other comfort items. This caused the power system to fail, and prompted work-around solutions such as the posting of signs on appliances offering warnings such as, "Don't run the microwave when the coffee pot is brewing." Eventually, this challenge prompted the purchase of larger generators and better distribution systems. Soon, refrigerators, freezers and air conditioning systems were added, causing additional outages to the system until the military eventually, through trial and error, identified the maximum power demands of the soldiers. Always supplying electricity for maximum demand resulted in highly inefficient power generation plans with too much power generating capacity. Generators were failing at high rates, since the majority of the time the generators were not running at the optimal load levels. A considerable amount of money was spent to correct this misallocation of electrical power resources. The same unmeasured pattern of supply and demand is occurring in our Command and Control (C2) systems today; however, instead of electrical power, bandwidth is the contested and misallocated resource.

During initial design, individual C2 systems are developed to support very specific tasks and missions, however, over time, as additional functionality and features are added, their bandwidth demand increases. Today, C2 systems are evolving and exchanging increasing amounts of data at a rate faster than tactical network bandwidth can grow. Additionally, as C2 system consolidation initiatives and cloud computing solutions are pushed to the tactical edge, even more traffic will be placed on the existing bandwidth. One answer is to buy more bandwidth; however, this approach runs into the same issues as the power problem described earlier in this paper: increased cost, complexity of implementation, hardware maintenance issues, and inefficient operations caused by building for peak demand resulting in excess capacity the majority of the time. Adding more bandwidth to support peak operations is not always possible for remote or small units on the forward edge of the battle space.

A key objective of Project MERCURY is to dynamically manage services in order to optimize their performance using available resources such as bandwidth to ensure continuous delivery of mission critical information. This requires managing the movement of data at the service level in order to ensure the commander's priority information is made available, while sustaining data exchanges needed to support future operations or sustainment activities. This solution requires dynamic enforcement of appropriate Service Level Agreements (SLAs) on individual services and groups of them, or service profiles, to ensure mission critical information is made available regardless of the available bandwidth and competing demands of other users of the C2 systems.

Although dynamic, intelligent and mission-focused allocation of all contested IT resources is in the purview of Project MERCURY; its first phase focuses on bandwidth allocation only. Moreover, it is important to distinguish between bandwidth allocation and bandwidth management. Bandwidth management is addressed by designated special programs and initiatives like WIN-T. Bandwidth allocation is an application-specific concern and needs to be addressed with an in-depth understanding and instrumentation of the services and applications utilizing the network, and ongoing analysis of the usage patterns dictated by each unit's missions.

3 BANDWIDTH ALLOCATION TODAY

Today, the commander cannot quickly align available bandwidth to the current mission needs because understanding of the bandwidth utilization is very fragmented and only available at a very technical level. Moreover, it is not possible to allocate available bandwidth to only specific capabilities needed for the current mission. Currently, C2 systems in the Tactical Operations Center (TOC) are designed to be operational 24/7 and support all phases of operations. However, Intelligence Surveillance and Reconnaissance (ISR) feeds are progressively demanding significantly more bandwidth due to the ability to provide full motion video in real time. Other sensors

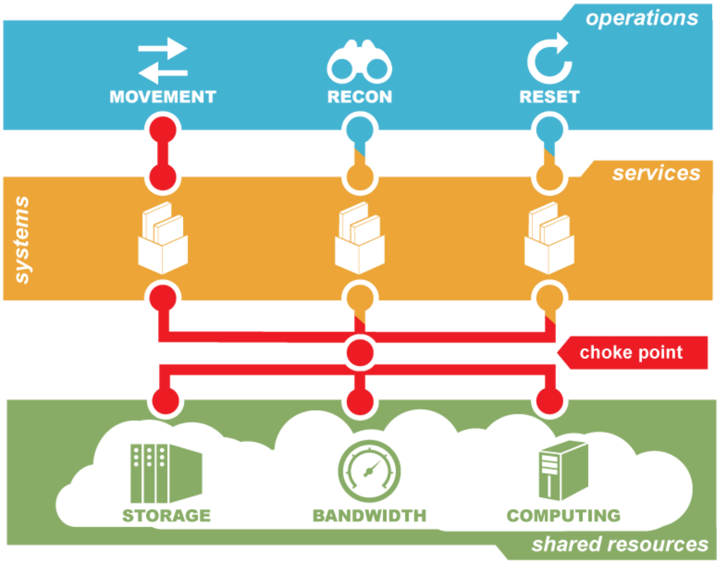

Figure 1: Operational Contention for Limited Resources

are being added to the commander's network to provide improved security and situational awareness information, also requiring more bandwidth. Figure 1 illustrates the bandwidth challenge that persists (in normal weather conditions) the majority of the time, even when minimal activity is occurring, as the soldiers and assets work to create the information needed to support the next operation. The challenge occurs when a potential target is identified, ISR assets are sent to track it, convoy and deployment systems become engaged to support blue force movement, security systems are engaged to protect the movement of assets, and additional data is pulled and shared globally in order to coordinate with theatre-wide assets and strategic leaders. Should weather degrade a communications signal or any other heavy demand be placed on the system, there is a significant probability that major data feeds will be interrupted. Should one of the key commander's data feeds be interrupted, the only on-the-ground, immediate corrective action is to literally "pull the plug" on another system deemed non-essential during the current mission, effectively cutting off an entire C2 sub-system from the network. This is the equivalent of unplugging the microwave in order to run the coffee pot; unfortunately we are talking about our warfighter's C2 systems.

Another major issue occurs when bandwidth is variable. During times of limited bandwidth or a short outage, network traffic gets severely backed up, and when the network is available again all essential and non-essential traffic floods the communication pipes. Currently, there is no effective way for the commander to tag information or systems by priority, and there is no way to adjust that priority as the mission changes.

The severity and effects of unplugging a C2 system for a period of time in order to support a critical mission are debatable. It is clear that a number of fundamental technical blind spots exist in the Army's understanding of how the C2 systems operate in a system-of-systems environment. Better understanding of these blind spots will enable the Army to better plan, manage and control the overall C2 system-of-systems performance in tactical environments with constantly changing bandwidth conditions.

The bandwidth allocation blind spots (areas where information is missing) are due to the lack of:

- Understanding of actual bandwidth utilization at the system, component and service level
- Standard technical metrics for tracking system data trends, patterns and usage
- Sharing of metrics on bandwidth availability and utilization across system boundaries
- Standardized machine readable SLAs
- Standardized taxonomy for prioritization of services
- Intuitive controls commanders can use to manage data transmission loads
- Dynamic automated service balancing based on changing bandwidth availability
- Dynamic controls for bandwidth consumption at the individual TOC level
- Insight into the data needs for different mission types and phases

Can these technical blind spots be overcome by adding bandwidth and continued reliance on the "pull the plug" approach to data management? Absolutely not, because this challenge is not limited to major TOC operations with multiple ISR feeds that consume tremendous amounts of data. According to a recent field report from the 101st Airborne Division (2nd Brigade Combat Team, 2011), this issue appears at the lowest level of the C2 system at the Battalion and Company TOCs. Since current Army doctrine is pushing C2 technology solutions down to the lowest level, it is obvious that having many blind spots in the architecture will make it very challenging to understand and predict data transportation requirements if information on actual system usage is limited or too coarse-grained. In the report from the 101st, not only did they need to provide bandwidth for the C2 systems they deployed with, they needed to do it using several different systems: Joint Network Node (JNN), Command Post Node (CPN), Secret Internet Protocol Router (SIPR) and Non-secure Internet Protocol Router (NIPR) Access Point (SNAP), Harris 7800W SSS, commercial line-of-sight (LOS) microwave radio, and Harris 117G radios. To make matters worse, the exchange server running at one of the sites ran for 12 hours and consumed 300 Megabytes of data but only sent two emails with no attachments. Clearly, there is an opportunity to improve bandwidth allocation and overall system-of-systems performance.

The Army, more specifically the U.S. Army Training and Doctrine Command (TRADOC), has recognized the need to better manage the flow of data to support commanders, and has documented this need in three places in the Mission Command Essential Capabilities document (TRADOC Capabilities Manager, 2011) as quoted below:

- A Robust Network … enables the timely flow of mission command information in accordance with the commander's priorities
- Execute Tactical Network Operations: Provides operational control of transport resources, in accordance with command priorities, to meet mission requirements … include the efficient and dynamic allocation of transport resources to maximize mission command application and service performance … Thus, commanders will have the capability to accomplish mission command … during disconnected, intermittent and low bandwidth conditions
- Command and Control On-the-Move: Provide commanders the ability to monitor the COP so as to maintain [Situational Awareness] SA and communications while away from the [Command Post] CP ... by allowing the commander to command and control a mission from the best location while making timely and informed decisions in response to changing conditions

With the move toward a Common Operating Environment (COE) and shared infrastructure, the "pull the plug" approach will cease to be effective as multiple services will be running on a physical server. While some units today manage bandwidth by IP address, this approach will also become ineffective because in cloud based deployments, various service requests originate from the same IP.

The need for intelligent, dynamic bandwidth utilization controls is real and it is recognized by TRADOC. With TOC consolidations in progress and emergence of COE, the next step is to identify a potential solution.

4 THE POTENTIAL SOLUTION

A solution to the bandwidth consumption problem can be developed today with existing technologies and emerging techniques. The ability to collect metrics on service usage, aggregate and average the service activities over time, display the information on a simplified dashboard and provide mechanisms to allow machine readable SLAs are all possible today. While some aspects of the solution were possible before with preclusively expensive software and hardware, today the solution can be created with existing government owned software, such as CERDEC CP&ID-developed Open Enterprise Service Management (OpenESM) and emerging WIN-T network prioritization standards and protocols. As C2 systems continue to adopt Service Oriented Architecture (SOA) and COE guidance, the conditions are being set to further enable the detailed collection of metrics at the service level.

Figure 2: Notional Allocation Interface

In order for the solution to be viable it will need to meet several operational requirements. It must provide the commander an easy to understand summary of the current and past behavior of the systems, and allow him to quickly prioritize allocation based on the current mission. To avoid having to train a specialist for each tactical unit, the propagation of this prioritization must happen automatically across all of the affected systems. To avoid huge costs in fielding too many TOCs, the solution needs to be software-based, seamlessly integrated with existing architecture. The allocation interface needs to be simple, with individual data services logically grouped to allow for efficient management during critical operations. The interface needs to be designed for use by the S6's across all levels of the Army and require minimal training. Figure 2 provides an illustration of a simple view needed to understand data use and resulting bandwidth consumption by profile.

Project MERCURY aims to identify the key characteristics required to support the local planning, management and control of data services in order to ensure that during a critical mission, the essential data is made available using the lowest possible bandwidth without manual shutting down of non-essential systems and services.

The characteristics of a successful solution are:

- Easy for commander to understand in stressful or unpredictable solutions
- Technology and system agnostic

- Automated system learning capability to improve data usage trends
- Provide fine grained control over individual or groups of services to enable the ability to adjust to situations in incremental steps (bandwidth, operating conditions, mission scope, participating systems)
- Can be integrated with existing systems and networks
- Operated by military personnel, no field service representative (FSR) supported required
- Provide real-time visibility into service usage and bandwidth consumption at the individual service level or in aggregate profile views

In a future scenario with the Project MERCURY capability employed, should there suddenly be a bandwidth issue, an S6 Non-Commanding Officer (NCO) monitoring the network could look at the MERCURY dashboard and immediately identify the individual services consuming the greatest amount of bandwidth. Once those services are identified (not possible today), the NCO can de-prioritize these services, pause the service calls for a time, or turn off the services indefinitely. This corrects the bandwidth utilization issue for the commander in just a matter of minutes. The NCO could also create a new rule for the offending services assigning them the appropriate resource limits should this particular mission be executed again in the future. Additionally, once the rules for the service have been established and tuned they can be provided to other units or even to TRADOC to help in the planning and development of data service Standard Operating Procedures (SOP) for mission command systems. Most importantly, in a matter of minutes, the NCO could rebalance the entire service portfolio to ensure the commander is receiving the information needed at the time it is needed, without shutting down the entire system-of-systems. In this scenario, rather than unplugging the coffee maker, the S6 sets the timer to have it start brewing after the microwave is finished cooking.

5 PROFILE-BASED BANDWIDTH UTILIZATION

In an effort to simplify bandwidth utilization while enabling the commander to make decisions based on empirical and historical data CERDEC CP&ID set out on the course of developing a profile-based bandwidth utilization solution – project MERCURY. The approach relies on the concept of a mission profile. A mission profile is a collection of related IT systems and services needed to support a single mission. For example, a logistics related mission may include an inventory service, a property book service, a part ordering service and a maintenance service. The collection of these specific services needed to support a logistics mission can be grouped into a mission profile. Mission profiles make managing groups of services much more efficient. Project MERCURY monitors and manages the groups of services in mission profiles using three major components:

- **Collection Agents** – agents that monitor bandwidth usage and make possible establishing realistic thresholds for each service
- **Allocation Agents** – agent that enforce the service level agreements which result based on the mission profiles, the service thresholds, available bandwidth and commander prioritization

- **Management Dashboard** - provides a rich, intuitive way to establish and manage the mission profiles, service bandwidth usage and view comprehensive historical metrics

5.1 COLLECTION AGENTS

Collection agents are low-level, embedded components that monitor and collect bandwidth usage across the utilized infrastructure, such as application and web servers, routers and operating systems. This information is later used to establish realistic thresholds for each service and help in composing mission profiles based on usage. The collected information includes metrics such as request and response time, payload size, frequency of requests and number of unique consumers. Several of these agents already exist for common infrastructure components as part of OpenESM project, but additional ones may need to be created to support the full range of Army systems and services. As part of project MERCURY, a simple interface will be published to make it possible for anyone to implement the required agent.

5.2 ALLOCATION AGENTS

Allocation agents enforce the allocated bandwidth utilization budget. They are the accountants of the mission profile-based approach. Allocation agents must be deployed across the infrastructure in application servers and network components like routers. These allocation agents integrate with bandwidth management methods that already exist through Army and DoD-based network management solutions like WIN-T and others. Allocation agents in project MERCURY make it possible to propagate the commander's intent and priority all the way down to individual packets of information flowing across the changing network conditions.

5.3 ALLOCATION DASHBOARD

The Allocation Dashboard, a web-based application, is a key component of project MERCURY and is made up of four modules. These modules, described in Figure 3, include the Commander's Dashboard Module, Service Metrics Module, Service Management Module and Profile Management Module. Sections below describe each module in detail using a simplified logistics mission use case. The screens are prototypes and are still undergoing major usability analysis and will change to address results of the analysis.

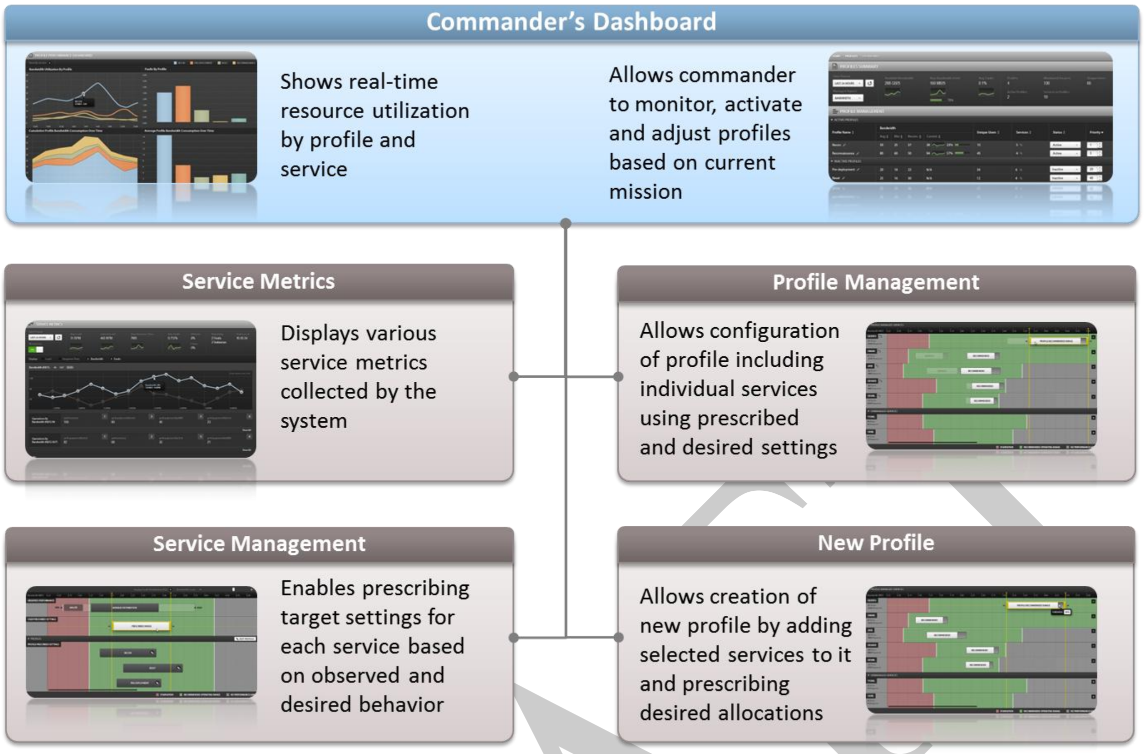

Figure 3: Allocation Dashboard Modules

5.3.1 Overall MERCURY Concept of Operation

Project MERCURY, upon its full realization, will enable optimal usage of available bandwidth. The following section describes what each of the involved participants in the typical workflow will contribute to the identification, configuration and run-time operation of various mission profiles that enable mission focused resource allocation. Figure 4 describes the notional sequence of operations enabled by various components of project MERCURY that results in configurations and activation of mission profiles based on the commander's current priority.

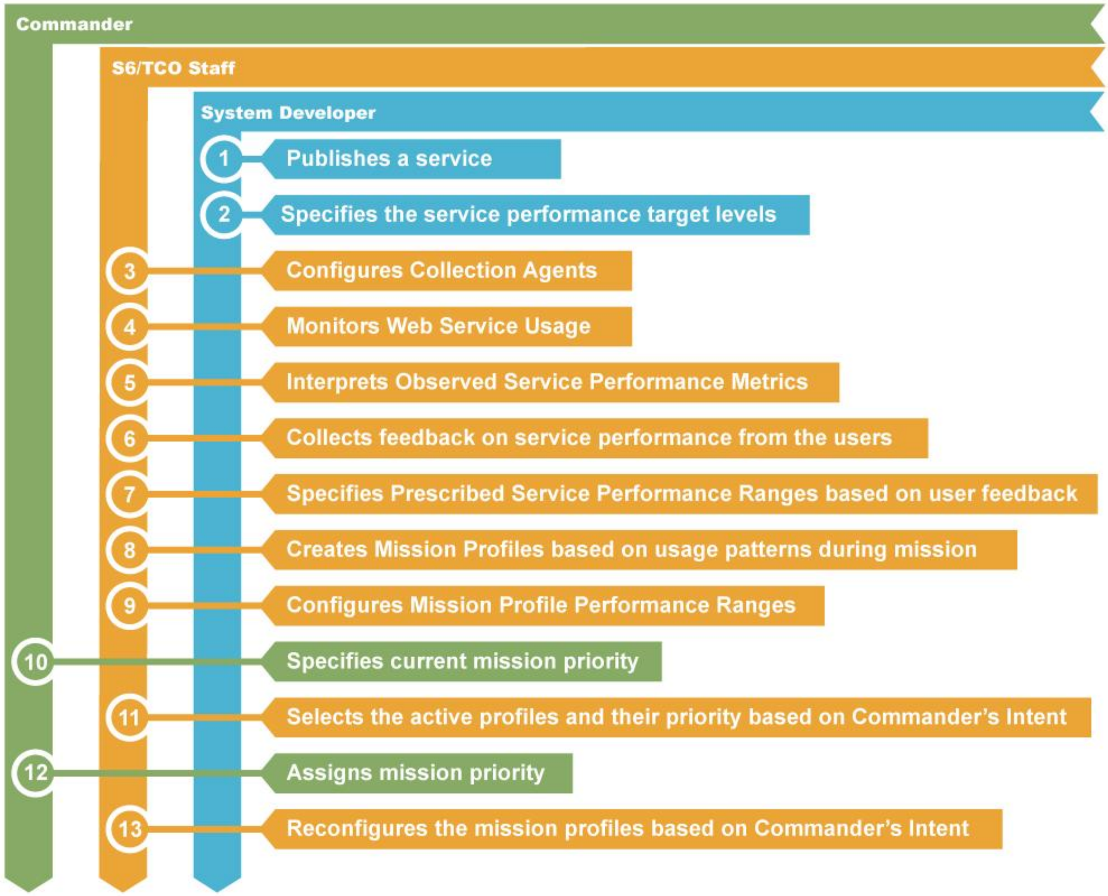

Figure 4: Project MERCURY CONOPS

5.3.2 Service Metrics Module

Low level service metrics collected by collection agents are displayed in the Service Metrics Module. This module helps the system administrators (S6) to examine the individual service behavior. The top section of the screen shows basic service metadata like service name, provider and known points of contact. This is information most often found in a service registry. The service metrics section provides access to metrics by configurable time period. In the example illustrated in Figure 5, metrics from the last 24 hours are displayed.

Figure 5: Service Metrics Module

Metrics include average load (requests per minute), current load and average response time. Each statistic is accompanied with a sparkline – a small, word-sized graphic, to indicate the most recent trend. For example, in the average response time metric, the user can see that there was a spike in the latter part of the selected period. The middle part of the Service Metrics Module contains a graph section. This section can display one or more of the user selected variable metrics such as load, response time, bandwidth, or service faults. Insight such as increased faults during decreased available bandwidth can be gained by overlaying multiple metrics.

5.3.3 Service Management Module

The Service Management Module illustrated in Figure 6 centers around two important concepts: *Observed Performance* and *User Prescribed Performance*. Observed performance when analyzed with regard to bandwidth, indicates how the service performs across the various bandwidth values. Collection agents enable analysis of each service by aggregating detailed metrics paired with the corresponding bandwidth condition.

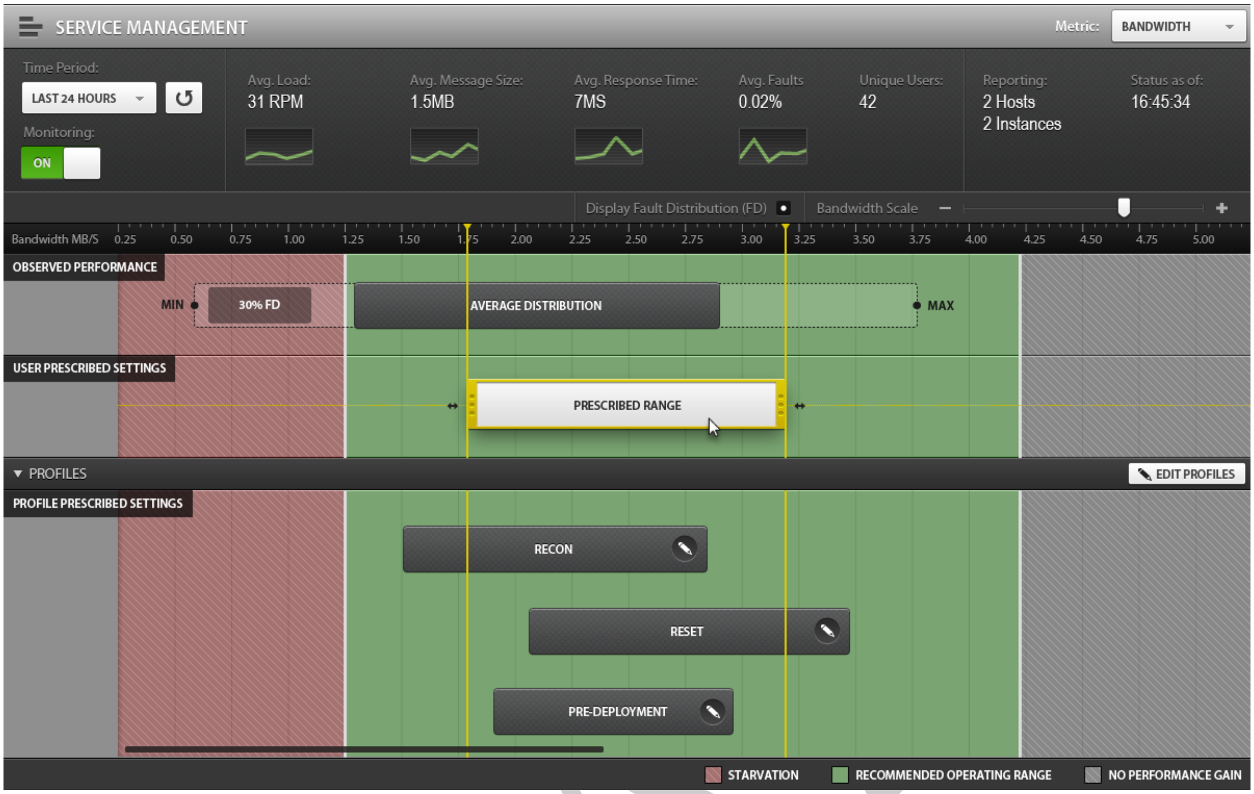

Figure 6: Service Management Module

Project MERCURY segments each service's observed performance into three zones: 1) Starvation Zone, 2) Recommended Operating Zone, and 3) No-Performance Gain Zone. Each service has a range where critical performance metrics such as faults become unacceptable. This range is known as the **Starvation Zone**. In Figure 6, the starvation zone is the leftmost red area. The goal of effective bandwidth allocation is to prevent services from entering the starvation zone. The green zone represents the optimal or **Recommended Performance Zone**, where the service responds in a timely manner and meets all user expectations – it performs according to the SLA. The last zone of the observed performance metric is the **No-Gain Zone**, indicated in grey. Most services will reach a point when allocating more bandwidth simply provides no additional benefit - thus reaching the No-Performance Gain Zone.

In order to establish meaningful and accurate observed zones, collection agents will need to instrument and collect services for a period of time. Once the zones are established, project MERCURY users can begin the process of prescribing non-default ranges, thus creating user prescribed settings. These settings allow the user, usually a technical expert familiar with the service operations, to set the desired ranges for each service while taking into account the observed performance metrics and the established zones.

The bottom section of the Service Management Module provides a listing of all mission profiles that this particular service participates in. By clicking on the specific profile, a user can access the Profile Management Module.

5.3.4 Profile Management Module

As mentioned throughout this paper, mission profiles are the novel part of project MERCURY's approach to dynamic bandwidth allocation because they enable more efficient management of groups of services. Each mission profile combines related services into a unit that can be managed together according to settings customized to each service in the profile. Since profiles are mission-based, a commander can prioritize and manage them without needing to understand the low level details. The management agents, together with project MERCURY infrastructure, take care of budgeting individual service's bandwidth usage according to the ranges previously set by the user based on observed performance metrics gathered by the collection agents. Figure 7 illustrates the interface for managing a mission profile.

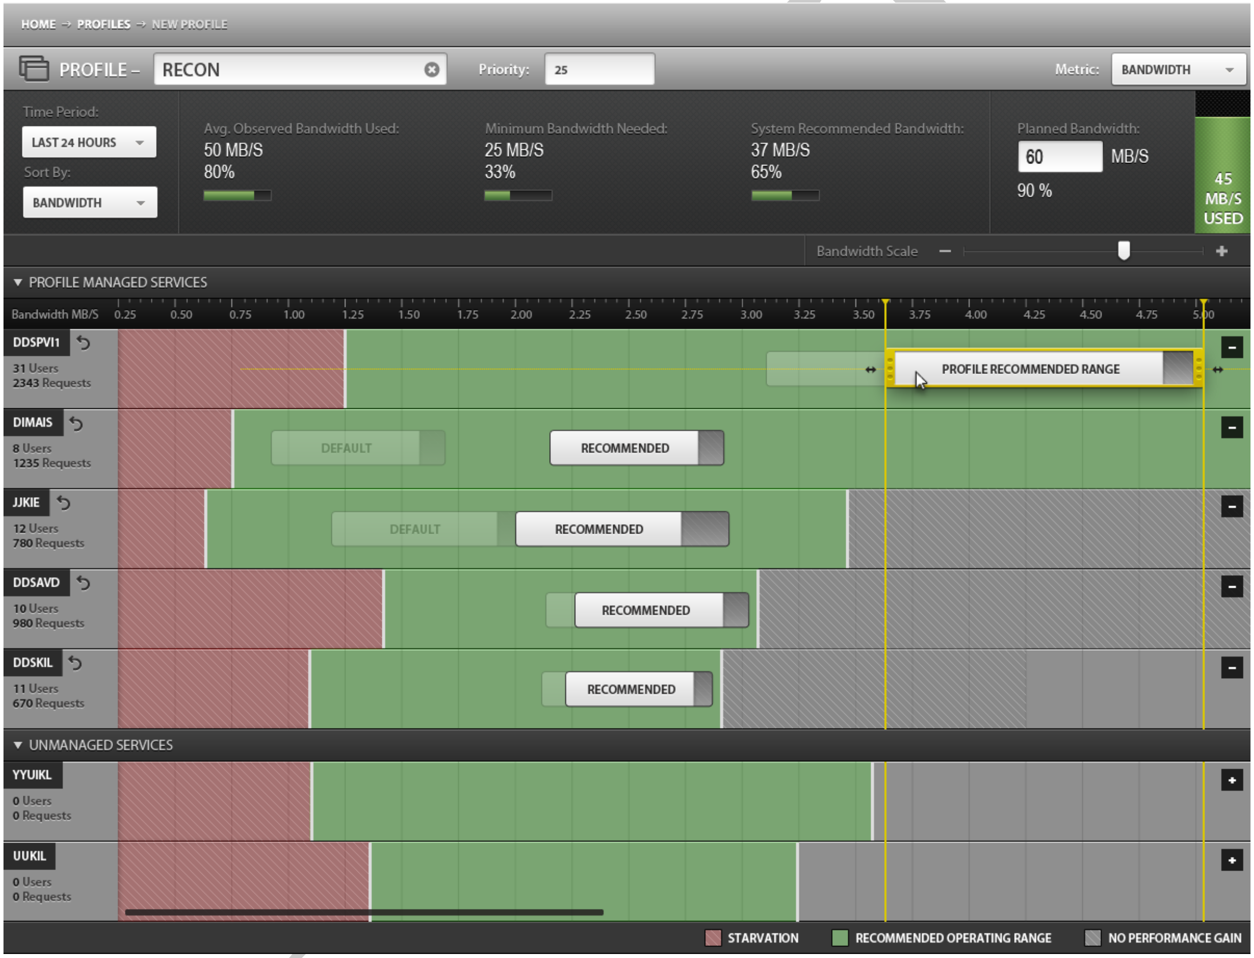

Figure 7: Profile Management Module

The top of the screen shows the current profile name and its priority. The priority ranges between one and 50. The value indicated in the profile management screen is only for planning purposes, it can be adjusted at run-time by the commander depending on the mission's priority. In the top left section, the user is given an option to select the period of time and the sorting metric. In this case, bandwidth over 24 hours is selected. The section to the right provides key metrics for managing the overall mission profile. Average observed bandwidth shows what the collection agents have observed about the combined bandwidth usage of the services that are currently part of the mission profile. The minimum bandwidth needed metric indicates the least value needed to prevent any of the services from going into their starvation zone. System-recommended bandwidth shows the average value that would keep the services of this profile performing in close to optimal range. The planned bandwidth value allows the user to plug in various values to see their effect on the overall mission profile performance. The green bar on the far top right provides a visual indicator of how much of the overall available bandwidth is utilized by the profile.

Figure 7 shows two groupings of services: profile managed services and unmanaged services. To add a service to the mission profile, the user can drag it from the unmanaged section. Future versions of project MERCURY will be able to suggest likely candidates for a profile based on the services already selected. Once a service is chosen to be part of the profile, its bandwidth usage affects the profile's allocation and actions performed on the profile, such as increasing its priority, will trickle down to the service. For example, if the overall allocation to the mission profile increases, the service may be given a larger allocation as well.

Three different ranges are visible for the profile managed services: default, recommended, and profile recommended. Default range, shown as a ghostly line, reminds the user where the service default recommended range is, regardless of what mission profile it is a part of. The recommended range indicates what the system thinks the service should be allotted based on the current mission profile and the available bandwidth value selected. The profile recommended range is a range assigned by the user to this service as part of this mission profile. For example, a user may choose to allocate more bandwidth to a UAV feed as part of a recon mission profile.

5.3.5 Commander's Dashboard Module

The commander's dashboard module is a culmination of the low level service metrics collection, and management and mission profile creation. It provides the commander a way to see only the essential information needed to manage the available bandwidth according to the current priority, while not starving the non-priority services. Figure 8 is a snapshot of the situation in which two mission profiles are active and two are inactive. The top section of the screen provides a familiar view of the available and used bandwidth as well as the average faults, all with sparklines to show trends over the selected period. The top right section has more top level statistics such as the number of monitored services and services that are currently part of mission profiles. It also shows the total number of known users.

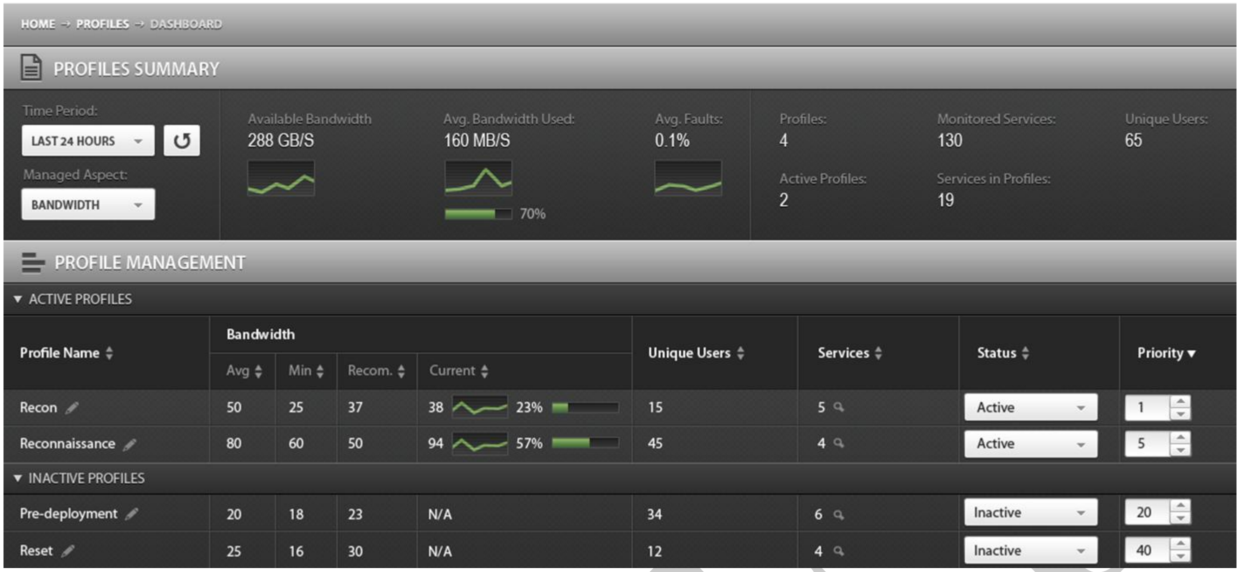

Figure 8: Commander's Dashboard - Profile Management

The Profile Management section of the Commander's Dashboard gives the commander a way to set a mission profile as active and to change its priority at run-time. It also provides key metrics for each profile, such as average bandwidth consumed, recommended bandwidth value and the current usage indicators. Unique users and number of managed services can also be useful when making changes to the profile.

Figure 9: Commander's Dashboard - Profile Performance Dashboard

Another screen that is part of the Commander's Dashboard is the Profile Performance Dashboard. It gives the highest level view of each profile's impact on the bandwidth utilization. The commander can immediately tell what profile is consuming the largest portion of the available bandwidth. In Figure 9 it is the blue area – Recon profile. The commander can also see when the profile's bandwidth usage spiked. This view provides the commander an actionable and clear roll up of a massive amount of data collected over the chosen period, across technology stacks, platforms, and operating systems, and it provides a way to adjust overall mission allocations or drill down to individual service limits.

6 CONCLUSION

Project MERCURY is moving ahead with development of this critical operational capability. It will be based on existing and emerging technologies developed by Army and industry for use with existing and future systems and services. It will be fully integrated at all levels of technology stack including network (WIN-T) and service management and monitoring capabilities provided by OpenESM. It will be aligned with the latest guidance from COE and Command Post Computing Environment in order to address current and emerging challenges and to ensure the fastest possible deployment to the soldiers. While bandwidth is the most critical choke point today, the mission profile approach outlined in this paper can be applied to managing other constrained and contested resources such as memory and processing power. As more and more systems and services collapse onto the same hardware resources, as directed by COE, effective management of these choke points will be critical.

7 REFERENCES

2nd Brigade Combat Team, 1. A. (2011). *Interconnecting Diverse Satellite and Line of Site Networks in Afghanistan.*

Army CIO/G6. (2010). *Army Software Transformation Strategy.* Retrieved from https://www.us.army.mil/suite/doc/22399356

TRADOC Capabilities Manager. (2011). *The Mission Command Essential Capabilities White Paper.* Ft Leavenworth, KS: TRADOC.

8 ACRONYM LIST

C2	Command and Control
C4ISR	Command, Control, Communications, and Computers Intelligence Surveillance and Reconnaissance
CERDEC	Communications-Electronic Research, Development and Engineering Center
COE	Command Open Environment
CP	Command Post
CP&ID	Command, Power and Integration Directorate
CPN	Command Post Node
FSR	Field Service Representative
ISR	Intelligence, Surveillance and Reconnaissance
IT	Information Technology
JNN	Joint Network Node
LOS	Line of Sight
MERCURY	Mission Engaged Research in Cloud Use for the Army
NCO	Non-Commanding Officer
NIPR	Non-Secure Internet Protocol Router
OpenESM	Open Enterprise Service Management
S6	System Administrator
SA	Situational Awareness
SIPR	Secret Internet Protocol Router
SLA	Service Level Agreement
SNAP	SIPR and NIPR Access Point
SOA	Service Oriented Agreement
SOP	Standard Operating Procedures
TOC	Tactical Operations Center
TRADOC	U.S. Army Training and Doctrine Command