Net-Enabled Command Capability

INCREMENT 1 PHYSICAL ARCHITECTURE

Version 1.0

31 October 2007

Prepared by:

Net-Enabled Command Capability
Joint Program Management Office (JPMO)
P.O. Box 4502
Arlington, VA 22204-4502

In Collaboration with the NECC Component Program Manager's for the Navy, Air Force, Army, Marine Corps, USJFCOM and DISA.

DISTRIBUTION – Distribution authorized to DoD and DoD Contractors only; for administrative/operational use (October 2007). Other requests for this document shall be referred to the NECC Program Management Office.

DESTRUCTION NOTICE – For unclassified, limited documents, destroy by any method that will prevent disclosure of contents or reconstruction of the document.

APPROVAL PAGE

Approved by: _____ **Date:** 10/31/07

MARK J. KUZMA
Chief Engineer,
Net-Enabled Command Capability (NECC)

EXECUTIVE SUMMARY

The physical architecture defines the server environment that will host NECC Capability Modules (CMs) for Increment 1. This environment uses a standard, virtualized server platform called a Global Information Grid (GIG) Computing Node (GCN). GCNs deploy at both local and enterprise computing sites. They also provide a uniform deployment and security environment for NECC CMs. The virtualized environment of the GCN requires that NECC bundle its software into virtual machine packages called Capability Packages (CPs). NECC provisions CMs by deploying these CPs onto GCNs.

REVISION HISTORY

REVISION NUMBER	REVIEWER / ORG	CHANGES	REVISION DATE	DATE ENTERED	NAME OF PERSON ENTERING CHANGE
1.0	Smiley/MITRE	Baselined version.	24 October 2007	24 October 2007	Mark Smiley
	IGJ/BIT NECC Documentation Support DLB/SPEC	Reformatted and provided document final go through. Reset document settings, figures and table definitions; removed the DRAFT watermark; generated a new TOC.	31 Oct 2007	31 Oct 07	Issa Jones Dustan Brown

TABLE OF CONTENTS

APPROVAL PAGE ..... **II**

EXECUTIVE SUMMARY ..... **II**

REVISION HISTORY ..... **III**

1 STRATEGY **5**

1.1 RELATED DOCUMENTS 5

2 KEY CONCEPTS **5**

2.1 VIRTUALIZATION 5

2.1.1 Hypervisors *6*

2.2 GIG COMPUTING NODES 7

2.3 OPERATIONAL NODES 9

3 IMPLEMENTATION **11**

3.1 GCN SPECIFICATION 11

3.2 OPERATIONAL NODES 12

3.2.1 Enterprise Sites *12*

3.2.2 Operational Sites *12*

3.3 COOP AND REDIRECTION 14

3.4 WAIVER PROCESS 15

4 CONCLUSION **15**

LIST OF FIGURES

Figure 1: Generic Virtualized Server Environment 6

Figure 2: GCN with Type 1 Hypervisor 7

Figure 3: GCN with Type 2 Hypervisor 7

Figure 4: GCN Hosting Multiple CPs 8

Figure 5: GCN Containing Multiple Servers 9

Figure 6: NECC Architecture Taxonomy 10

Figure 7: CP Deployed on a GCN 11

Figure 8: Operational Site Class Hierarchy 14

LIST OF TABLES

Table 1: GCN Specification 11

Table 2: Terms of Use Metrics 12

Table 3: Classes of Operational Site 13

Table 4: Remotely Supported Operational Sites 14

1 STRATEGY

In Increment 1, Net-Enabled Command Capability (NECC) will provision Capability Modules (CMs) onto standardized hardware called Global Information Grid (GIG) Computing Nodes (GCNs). GCNs provide a common, Joint computing environment based on cutting-edge virtualization technology developed by industry. High-capacity, data-center class Enterprise Sites host Enterprise GCNs that can scale to large numbers of users. Operational Sites host Local GCNs that support the capability needs of a smaller set of users, for example, the staff at a tactical headquarters. Operational sites often have networks characterized by Disconnected, Intermittent, or Limited (DIL) communications. NECC Capability Packages (CPs) that implement a CM will be engineered to take these network limitations into account.

This GCN-based architecture ensures that software deployment and operating costs are lower, the deployment footprint shrinks, and capabilities can be deployed more rapidly. Additionally, GCNs provide a uniform accreditation environment, which improves overall security and simplifies the security process.

1.1 Related Documents

- NECC Increment 1 Software Architecture, version 1.0
- NECC Increment 1 Data Architecture, version 1.0
- NECC Increment 1 Technical Operations Architecture, version 1.0
- NECC Increment 1 Information Assurance Architecture, version 1.0

2 KEY CONCEPTS

2.1 Virtualization

A "virtual machine" is a tightly isolated software container that can run its own operating systems and applications as if it were a physical computer. A virtual machine behaves exactly like a physical computer and contains its own virtual (i.e., software-based) Central Processing Unit (CPU), Random Access Memory (RAM), hard disk and Network Interface Card (NIC).

"An operating system can't tell the difference between a virtual machine and a physical machine, nor can applications or other computers on a network. Even the virtual machine thinks it is a 'real' computer. Nevertheless, a virtual machine is composed entirely of software and contains no hardware components whatsoever. As a result, virtual machines offer a number of distinct advantages over physical hardware."[1]

Virtual machines possess four key benefits:

- Compatibility: virtual machines are compatible with all standard x86 computers.
- Isolation: virtual machines are isolated from each other, as if physically separated. Although if the physical machine on which they are based crashes, naturally all the VMs hosted on it will crash, too.

[1] http://www.vmware.com/overview/vmachine.html

- Encapsulation: virtual machines encapsulate a complete computing environment.
- Hardware independence: virtual machines run independently of underlying hardware.[2]

The move towards virtualization grew out of a need by IT departments for a simplified deployment environment, and for a more efficient use of available server capacity. It achieves these goals by adding a virtualization layer on top of the underlying platform (see Figure 1). This added layer abstracts the hardware and Operating System (OS), thereby providing a uniform and simpler deployment environment. Furthermore, since the virtualization layer can host multiple virtual machines, it becomes possible to maximize a server's available resources. For example, relational databases often use only a fraction of a server's CPU, which then sits idle. With virtualization, it is possible to add applications that use that formerly idle CPU, while maintaining true application isolation.

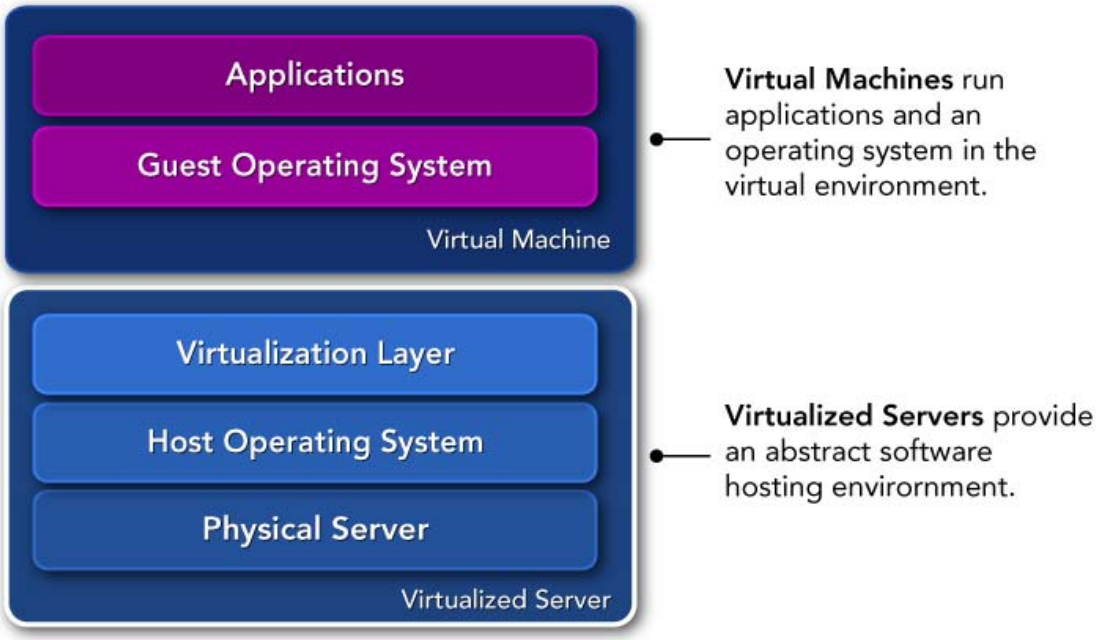

Figure 1: Generic Virtualized Server Environment

2.1.1 Hypervisors

A hypervisor is a virtualization platform that allows multiple operating systems to run on a host computer at the same time. The term usually refers to an implementation using full virtualization. Hypervisors are currently classified in two types.[3]

A Type 1 hypervisor (or Bare-Metal Architecture) is software that runs directly on a given hardware platform.[4] VMware's ESX Server is an example of this. In some cases, it is possible to embed the hypervisor in the server's Read-Only Memory (ROM) (see Figure 2).

---

[2] Ibid.

[3] http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/eicay/eicay.pdf

[4] http://en.wikipedia.org/wiki/Hypervisor

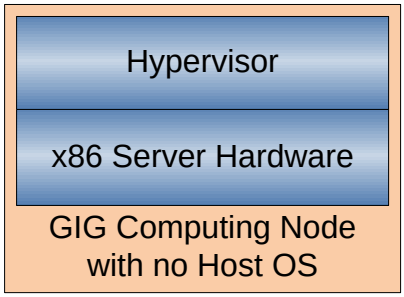

Figure 2: GCN with Type 1 Hypervisor

A Type 2 hypervisor (or Hosted Architecture) is software that runs within a host operating system[5] (see Figure 3).

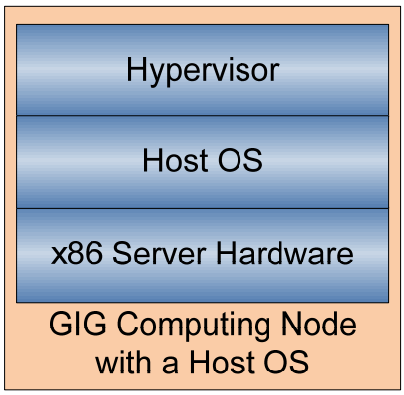

Figure 3: GCN with Type 2 Hypervisor

2.2 GIG Computing Nodes

A GCN is a common Joint computing environment, based on commercial virtualization technology. A GCN consists of a collection of one or more physical servers, each of which contains a hypervisor, such as VMWare's ESX Server (see Figure 2). If the hypervisor is of type 2, each GCN server also includes a host operating system (see Figure 3).

GCNs bring the benefits of virtualization to the Department of Defense (DoD). They establish a common, Joint virtualized server environment into which programs throughout DoD can provision their software (see Figure 4). They also have the potential to function as a universal security and accreditation environment[6]. NECC has adopted the GCN as its target deployment environment.

There are two main types of GCNs, based on their target operational node: Enterprise GCNs and Local GCNs. Enterprise GCNs reside at enterprise sites and can service any user across the enterprise. They maintain state. Within a CM, the CPs deployed to enterprise sites must meet

[5] Ibid.

[6] Refer to the NECC Increment 1 IA Architecture for more information.

capacity and performance numbers sufficient for the entire GIG. They have the responsibility to maintain the global state of the CM across the enterprise to support Continuity of Operations (COOP).[7] An example of this would be the full global blue force track picture.

Local GCNs are intended for operational sites and service local users at a local site. They provide onsite services when in DIL mode.[8] The CPs deployed at Local GCNs serve only the capability needs of Warfighters on the local network.

Locally deployed CPs can also make use of Enterprise CPs. This potentially complex relationship among Local and Enterprise CPs, however, is hidden inside the CM black-box which appears to the Warfighter as a single, coordinated set of services.

NECC is a software program that will provision software on to Service provided hardware. The Services have existing infrastructure programs, such as: Consolidated Afloat Networks and Enterprise Services (CANES) (Navy), and Battle Command and Control System (BCCS) (Army). These Service infrastructure programs are providing computing platforms that follow the GCN standard, too.

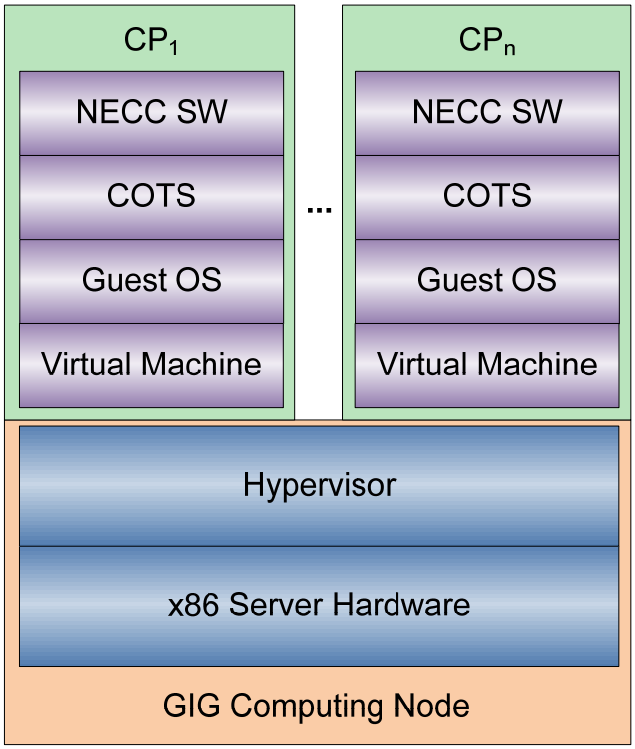

Figure 4: GCN Hosting Multiple CPs

[7] NECC Architecture Overview-v2-txw-30Jul07.ppt, by Dr. Mark Matthews.

[8] Ibid.

Figure 5 below depicts a GCN that contains multiple physical servers.

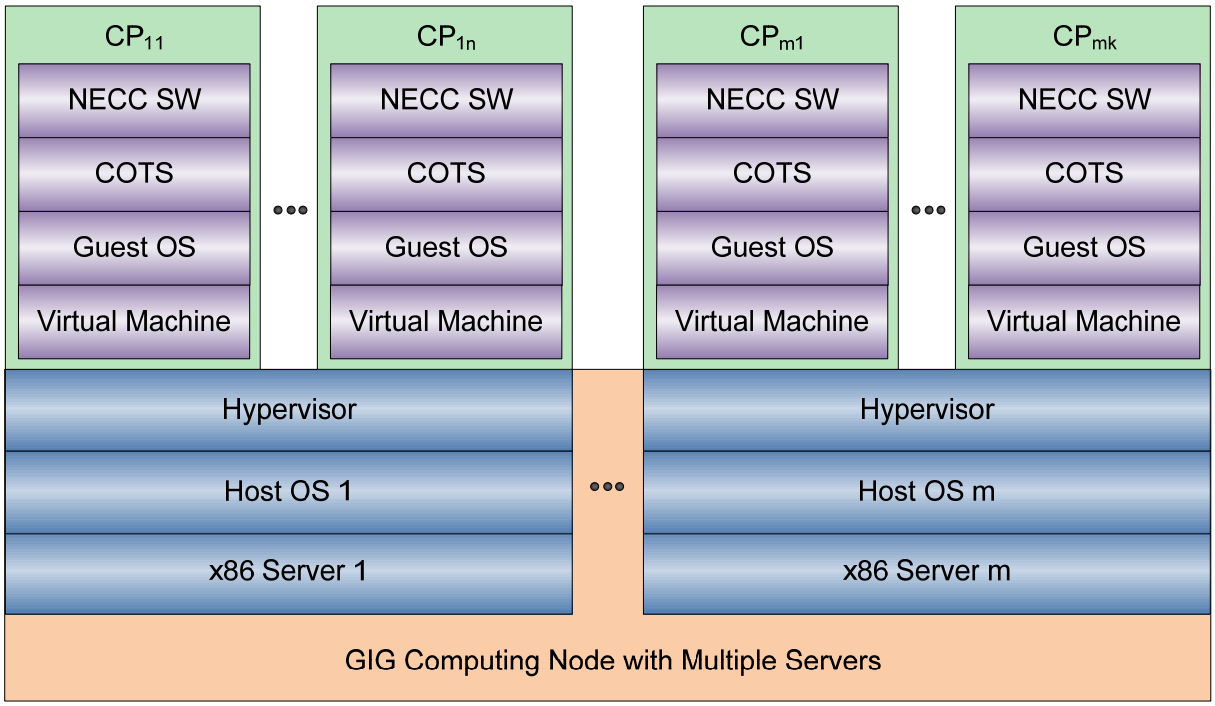

Figure 5: GCN Containing Multiple Servers

2.3 Operational Nodes

NECC divides operational nodes into two basic categories: enterprise sites and operational sites. Enterprise sites are data-center class facilities. They are fixed sites with generally high-bandwidth network communications, and large server capacity. A Defense Information Systems Agency (DISA) Enterprise Computing Center (DECC) is an example of an enterprise site. Operational sites are primarily military headquarters elements such as an Air Force Air Operations Center (AOC), an Army brigade Tactical Operations Center (TOC), or a Navy ship. These sites are usually mobile and often have DIL constraints. Figure 6 shows the relationship among these and other elements.

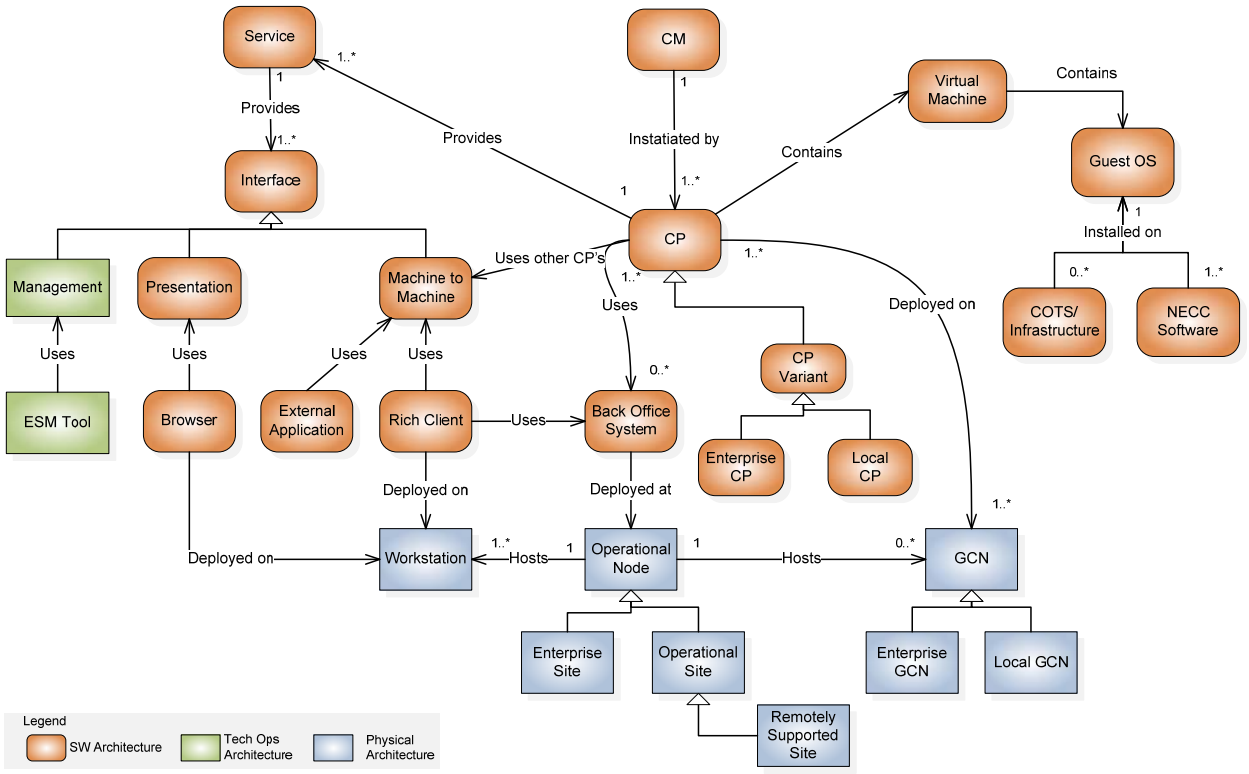

Figure 6: NECC Architecture Taxonomy

Operational sites that do not have NECC capabilities deployed locally are called remotely supported. Warfighters at these locations have full access to NECC capabilities using reach-back to enterprise sites. Here too, the Warfighter is shielded from most of the technical complexity since the CM presents a single interface point. For example, just as with commercial web sites, a Warfighter at a remotely supported site need only enter a single URL into a browser and will be transparently redirected to the most appropriate CP.

During NECC Increment 1, existing back-office systems will continue to play a key role. Since some NECC services will rely on these systems, especially during the initial roll-out of Increment 1 capabilities, and because some of them will not yet have made the transition to the new GCN environment, NECC may initially field certain back-office systems on non-GCN compliant hardware. In keeping with the Service Oriented Architecture (SOA) adopted by NECC, these deployments will usually be to a very small number of Enterprise Sites. Operational sites may continue to operate these systems, provided by the Programs of Record (PORs), during the transition.

Operational Nodes also host workstations, which are computers that operators interface with. Workstations may run Rich Clients, such as Command and Control for the PC (C2PC).

3 IMPLEMENTATION

3.1 GCN Specification

Local and Enterprise GCNs implement the same virtual environment. It uses the industry-leading, hypervisor product VMWare Infrastructure Enterprise. Viewed as a set of logical layers, physical servers lie at the bottom of the stack. All physical servers must use the x86 CPU. A GCN can contain more than one physical server (see Figure 5).

VMWare's ESX Server product occupies the layer above the physical server. It acts as both the host OS and the virtualization layer. CPs, reside above the hypervisor layer. They run inside the virtualized environment provided by a hypervisor, such as ESX Server. This is depicted in Figure 7.

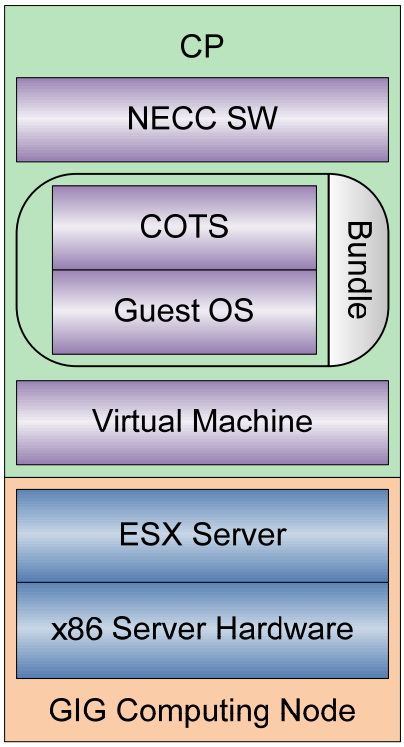

Figure 7: CP Deployed on a GCN

Table 1 lists specific products and hardware versions for the GCN.

Table 1: GCN Specification

Item	Version
Hypervisor	VMWare Infrastructure Enterprise version 3
Host OS	VMWare ESX Server. Part of VMWare Infrastructure Enterprise version 3.
Hardware	X86 CPU based servers only. SPARC and PowerPC are not supported by the GCN virtual environment.

3.2 Operational Nodes

3.2.1 Enterprise Sites

Over the course of Increment 1, NECC plans to install software at approximately four geographically distributed, enterprise computing sites. These sites will contain Enterprise GCNs and will host Enterprise CPs in a data-center class environment. As part of its migration strategy, NECC may also deploy legacy Sun SPARC servers at enterprise sites in order to support back-office GCCS applications needed by NECC during their transition to the GCN environment.

3.2.2 Operational Sites

Ensuring that local servers meet the GCN specification is a Component Program Management Office (CPMO) responsibility. Local GCNs may contain one or more x86 servers with sufficient resources to host VMWare Infrastructure and the requisite number of CPs. In order to help CPMOs with capacity planning, each CP will contain a "terms of use" that specifies the minimum resource requirements (see Table 2) needed to host and run it. These terms assume a virtualized GCN environment.

Table 2: Terms of Use Metrics

Resource	Description
Concurrent User Rating	How many concurrent users does this Terms of Use support?
System RAM	Minimum RAM that must be available to this CP.
Storage	Hard-drive space need for the application and its data files. Also includes external, e.g., SAN, storage requirements.
CPU	Minimum speed and quantity of processor.
Video	Any special graphics card requirements.
Networking	LAN and WAN bandwidth requirements.
Dependencies	NECC CM, NCES/Component services, or back-office system dependencies.

Operational site environments vary widely across the Components and even within a Component. They also depend greatly on the theater of operations during an on-going conflict. Understanding these environments is nevertheless crucial to planning an effective Joint C2 architecture. To make understanding and documenting these environments a tractable problem, NECC has abstracted them into classes of sites with similar environmental characteristics. These characteristics impact the technical solution; such as network connectivity and organic systems administration support. The list of these classes is shown in Table 3 and a simplified view of their connectivity is shown in Figure 8. These environment descriptions will help shape the Increment 1 design artifacts.

Table 3: Classes of Operational Site

Node Name	Component	Associated Military Echelons
JFACC	Air Force	Air component
AFFOR	Air Force	AFFOR - Operational, theater.
MAJCOM	Air Force	MAJCOM - Operational, resource provider
AF	Air Force	Numbered Air Force
Wing	Air Force	Wing - Tactical. Home station or deployed.
EAC	Army	Echelons above Corps
Corps	Army	Strategic land component operational headquarters at theater level
Division	Army	Command and control of assigned Army forces at tactical and operational levels
Modular BCT	Army	BCTs are stand-alone combined arms organizations: heavy brigade combat team (HBCT), infantry brigade combat team (IBCT), Stryker brigade combat team (SBCT), and Future Combat Systems (FCS) brigade combat team (FBCT)
Functional Brigades	Army	Specialized brigades assigned to both Army divisions and corps situationally dependent: aviation, maneuver enhancement, fires, RSTA, and sustainment
NMCC/ANMCC	Joint	OSD, Joint Staff, Service HQs
COCOM	Joint	Operational
Small MAGTF Afloat	Marine Corps	MEB CE MEU CE (JTF Enablers)
Small MAGTF Ashore	Marine Corps	MEB CE MEU CE (JTF Enablers)
MEF (Large MAGTF)	Marine Corps	MEF CE (Potential JTF HQ or Functional Component HQ)
MSC	Marine Corps	DIV, Wing (TACC), LCE (CSSOC) (Potential Functional Component HQ)
MARFOR	Marine Corps	Operational
Force	Navy	Operational
Group	Navy	JTF
Unit	Navy	Command
Disembarked	Navy	Tactical

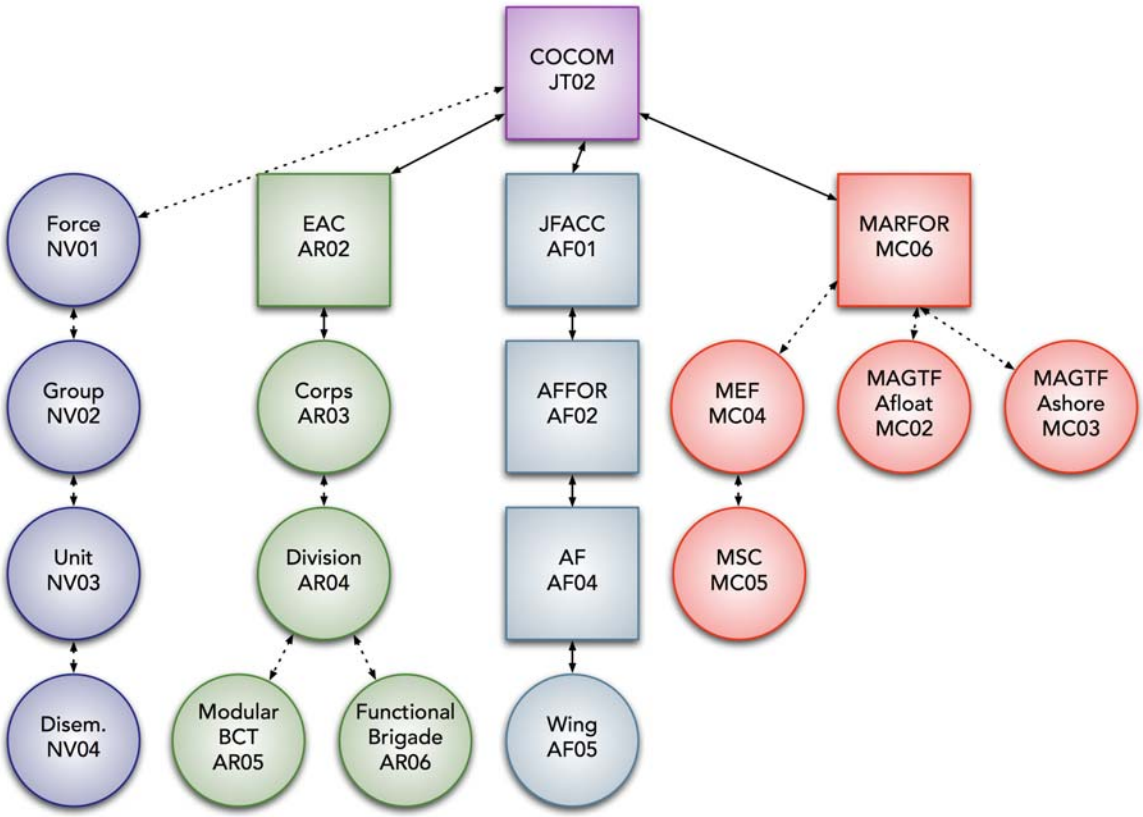

Figure 8: Operational Site Class Hierarchy

In the above figure, a dotted line indicates a wireless (radio/satellite) connection.

Table 4 lists the operational sites that have been designated as remotely supported sites. These sites will not host NECC capabilities locally.

Table 4: Remotely Supported Operational Sites

Component	Echelon
Air Force	MAJCOM
Air Force	Wing
Army	EAC
DISA	NMCC
DISA	COCOM
Marine Corps	MARFOR

3.3 COOP and Redirection

NECC enables Continuity of Operations for its capabilities by physically distributing CPs to GCNs around the world. Enterprise variant CPs maintain a copy of the global state of the data

managed by the CM. These CPs replicate and synchronize this data between themselves. If one of these CPs or sites goes down, the capability will continue to function using one of the remaining CPs.

If a Local GCN hosting an NECC capability stops functioning, clients of that capability can use an alternate Local GCN (perhaps one at a parent organization) or they can access it from an Enterprise GCN. COOP is managed internally by the components, which compose a CM. COOP does not cross CM boundaries. Materiel developers must implement the appropriate level of COOP support in their CMs taking into account any Local or Enterprise CP variants.

Redirection, transparently routing a service request to a CP, fulfills two key functions in the NECC architecture. First and most importantly, it allows the CM to offer a single, logical interface point. Clients, for instance, can access a service using a single URL and be redirected to an implementing CP. The secondary benefit is that this same mechanism gives the CM a COOP capability. Redirection can route the request to a functioning GCN if another is down. This fail-over is transparent to the client.

Unless they require a CM-specific redirection capability, materiel developers do not need to modify their CMs to use redirection. It functions as a layer above the CM software. Developers must, however, build a global data state replication mechanism into their enterprise variant CPs.

3.4 Waiver Process

Under certain circumstances, materiel developers may request and the Joint Program Management Office (JPMO) may grant waivers for compliance with the specifications contained in this document. Some specifications and standards can be waived, some cannot be waived, and others can only be waived for a limited period of time.

Specific waiver guidance for the Physical Architecture is as follows:

- Compliance of local servers to the GCN specification cannot be waived.
- The requirement for software intended for deployment on a Local GCN to be packaged as a CP cannot be waived.
- CPs intended solely for enterprise sites may receive a temporary waiver allowing their deployment outside of a virtualized (GCN) environment if there is a strong technical justification and the developer provides a transition plan and schedule to migrate towards full GCN compatibility.

4 CONCLUSION

This document describes the major NECC Physical Architecture concepts and implementation guidance applicable for NECC Increment 1. Virtual Machines, Hypervisors, GIG Computing Nodes, Operational Nodes, and other elements of the NECC Physical Architecture are defined and discussed in this document.