A Guide to Transition Your CommCare Infrastructure

A step-by-step guide to support organizations in transitioning their cloud-hosted CommCare applications to self-hosted and administered servers.

Cambridge, Massachusetts, USA | Cape Town, South Africa | New Delhi, India
www.dimagi.com

Table of Contents

Overview

The open source digital platform CommCare enables frontline organizations to build their own custom, permanent solutions on a professionally managed foundation - at any scale, anywhere. CommCare was originally developed by Dimagi, and today has evolved into being one of the most widely deployed Global Goods in existence.

CommCare HQ is a sophisticated, distributed software application, made up of dozens of processes and several pieces of third-party open source database software. It has been built for performance, reliability, security, and scale rather than simplicity. The vast majority of CommCare deployments use the global CommCare Cloud hosted instances of CommCare HQ via Dimagi's Software-as-a-Service (SaaS) model. This approach removes the need for the project to manage its own CommCare instance and also guarantees service availability and access to the latest platform features and security compliance with HIPAA and SOC-2. You can read more about the security and privacy standards offered by CommCare as a cloud-hosted solution here.

Thousands of governments and organizations run their CommCare applications as a cloud-hosted solution. While there are definitive benefits to hosting CommCare via the cloud, certain organizations or governments may opt to run CommCare on their own self-hosted or administered server. This may be due to a variety of reasons, including an organization's policies or regulations. Numerous organizations and governments have successfully transitioned their CommCare applications from being hosted by Dimagi's SaaS platform to an independently hosted instance of CommCare, without any required support from Dimagi.

> Dimagi has developed this guide to support interested organizations in transitioning their cloud-hosted CommCare applications to self-hosted and administered servers. This is written specifically for organizations who are either directly overseeing the transition between environments or are supporting a third-party (often a government).
> Numerous organizations have followed these steps to independently transition their CommCare applications without requiring any support from Dimagi.

As you get started, please feel free to refer to Dimagi's documentation, including CommCare Deployment Read the Docs and the CommCare Cloud GitHub documentation.

Phase 1: System Design and Planning

1.1 System Assessment and Design

The first step towards setting up a new instance of CommCare is to understand the intended long-term configuration of the system. This process includes:

- A full inventory of all planned environments (development, testing and production environments)
- An inventory of all mobile applications which will be hosted in the local environment
- An inventory of systems with which the CommCare instance will be integrated (for example, DHIS2, an Electronic Medical Record System, or a Data Warehouse).
- In addition, if users are currently collecting data in any of the applications, a migration plan to move any existing data from the SaaS hosted environment to the local instance must be designed.
- Projections on future addition of new users or new applications

1.2 System Sizing

Based on the planned number of users on the system, the program identifies the system configuration which is most suitable for the project. The table below summarizes default configuration sizes that can be used, though every application will have its own performance characteristics which can be further optimized.

Server Configuration	Scalability (Users)	Description
Single Server	< 1500	A single server on which all the pieces of the CommCare HQ software suite are installed
Micro Cluster	< 3500	Two servers running in parallel which provides higher capacity than a single server as well as better redundancy characteristics.
Small Cluster	< 15,000	A five-server cluster, which the parts of the CommCare HQ suite are distributed across
Large Cluster	> 15,000	A cluster of more than 5 servers, which the parts of the CommCare HQ software suite are distributed across.

CommCare Resource Model

In addition, the commcare-resource model is an open source, command-line tool designed to help a project estimate the infrastructure requirements for a project. Instructions for installing, configuring, and running commcare_resource_model can be found here. In order to run the resource model, users will be required to supply the following information:

- The number of CommCare users, and projected numbers, if the project is planning to scale, and planned dates when the users will become active.
- The existing total number of form submissions.
- The expected number of forms that an average user submits per month.
- New cases per user per month, and updated cases per user per month.

The output of this tool is a CSV file with a summary for each service, including the number of virtual machines and their resource requirements. It will also include detailed break-downs for each service on separate worksheets.

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely need to have the following information

- □ A link to the system architecture diagram including all planned environments and system integrations (sample system architecture diagram 1, sample system architecture diagram 2)
- □ Intended scale of the system documented, including a timeline to achieve this
- □ Planned server configuration documentation
- □ A list of all mobile applications which will be hosted in the local environment, including current released version and number of mobile users
- □ An answer to the question: “Are users currently collecting data in any of the applications listed above? If yes, how long have the applications been live?”

Checkpoint: Do not proceed to the next phase until the following have been completed

- System architecture diagram has been created including all planned environments and any systems with which CommCare will be integrated.
- Inventory of all mobile applications which will be hosted in the local environment, including current released version and number of mobile users has been documented.
- Intended scale of the system and timeline to achieve this has been documented.
- Server configuration has been finalized based on intended system scale and implementation phases. A sample system configuration diagram can be found here.

Phase 2: Hardware Procurement & Personnel Resourcing

As with any technology transition of this nature, all deployments of CommCare server or server cluster require planning, procurement of specific hardware, and skilled personnel to maintain.

2.2 Data Center Setup

The program identifies appropriate infrastructure to host the cluster. The program typically has three options for hosting, which are outlined below. The information below also includes key considerations as decisions are made regarding the data center.

Consideration 1: Physical Security

The physical security of hardware assets must be assured at all times. Compliance certifications (ISO 27001, SOC-2,CMMC) and procedures for the data center should be documented and audited prior to installing software on physical hardware. Note that Dimagi or other contract stakeholders may have contracted obligations for data security that will need to be addressed to ensure compatibility with the post-transition security posture. If necessary, Dimagi will provide the partner with any elements of compliance relevant to their contracted role with the future state of the system.

Consideration 2: Hardware

The number and type of servers required differs per cluster configuration; detailed information regarding hardware specifications may be found on the deployment documentation site, Hardware Resource Requirements for CommCare HQ. Investing in high quality Solid State Drives (SSDs) is critical to ensure that users don't experience system slowness and reduces the potential for data corruption.

- **Option 1: Enterprise Cloud Service Provider:** Enterprise Cloud Providers (Amazon Web Services, Microsoft Azure, Google Cloud Platform, etc.) provide public or managed private clouds, and offer on-demand cloud computing components. Users may create an account and declare the location of provisioned resources in the cloud management console.

- **Option 2: Local Cloud Service Provider:** A local cloud service provider is a vendor which provides physical infrastructure, software and full lifecycle management of the hardware infrastructure on which CommCare will be deployed. Dimagi has created a sample CommCare Hosting Vendor RFP for partners to extend when selecting a local CSP; an agreement in place should include Service Level Agreements (SLAs) which outline expectations for what is needed by the bidder for the architecture requirements of CommCare to be met

- **Option 3: On-Premise Data Center Setup:** Creating an on-premise data center involves procuring physical goods (servers, enterprise backup tools, high-speed internet connectivity), implementing maintenance and support plans for physical goods, acquiring security and compliance certifications, and preparing for physical security of the data center.

Consideration 3: Virtualization

Where possible, virtual servers hosted by a Cloud Service Provider (CSP) should always be preferred over physical servers. Maintaining and scaling the cluster is much less complex when working with a CSP; it is also generally more cost effective as the utilization of the resources can be optimized and the burden of hardware management is removed.

Adding virtualization to physical servers makes it possible to utilize the physical resources better, and makes certain maintenance tasks simpler. Installing and managing a virtualization layer requires highly skilled personnel.

Consideration 4: Network Connection

One of the biggest difficulties in maintaining physical server hardware is that to be reliable it must always be connected to the internet over a stable and reliable network. At a minimum, Dimagi suggests 940/880 Mbps speeds for any server which is to be used to deploy CommCare. Speed tests may be used to assess the connectivity in the location of the on-premise servers. Higher bandwidth may be required for larger scale projects.

Consideration 5: Power Source

A reliable power source is a critical component of managing physical hardware. Even with a generator, unless there's a sophisticated battery backup system in place, it will take a number of seconds for the generator to kick in, and even the slightest blip in the power supply will cause a server to shut off without warning. Thus, a well-planned system must be in place for maintaining consistent power even through a grid outage. Without this, users will experience downtime and the system will be at risk for data corruption as a result of powering off without warning. This is typically standard in all both global and local CSPs, but should be considered critical for on-premise setups as recovering from power failure corruption can cost much more in labor than the price of proper power redundancy.

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely need to have the following information

- □ A link to the server(s) specifications which will be used to host CommCare, including number of CPUs, amount of RAM the server / cluster has, and size / type (SSD or HDD) of the disk(s)
- □ A known understanding if CommCare will be hosted directly on these servers, or hosted within a virtual machine(s) within the hardware
- □ An established plan for expanding Disk space as needed based on growth
- □ Name of Internet Service Provider (ISP), the Mbps max burst download and upload speed, latency and bandwidth limits
- □ Link to the support plan and SLAs for when network issues arise
- □ Known understanding if the local subnet where CommCare will be installed is able to reach network locations outside of the local network
- □ Known understanding for how many minutes the Uninterruptible Power Supply (UPS) provides when an outage occurs
- □ Knowledge of how the UPS failover to safe shutdown is configured, and how long the safe shutdown takes to complete once power is lost.
- □ List of the compliance certifications (ISO 27001, SOC-2, CMMC) or procedures being followed at the physical data center, and auditing schedule for each
- □ Outline of physical security measures which are in place at the data center
- □ List all security appliances (web application firewall, intrusion detection monitoring, etc) which will be deployed along with CommCare services by type and version.
- □ Plan for ongoing system maintenance activities

Checkpoint: Do not proceed to the next phase until the following have been completed

- The program has signed an agreement with an Enterprise / Local Cloud Service Provider OR has completed setup of an On-Premise Data Center

2.3 Personnel Resourcing

The program identifies a team which will act as the Managed Service Provider (MSP), administering the CommCare HQ software suite installed in the identified infrastructure. Managing a deployment of CommCare requires a team of 1-3 part-time DevOps engineers, depending on the size of the deployment. A full outline of the MSP's responsibilities are included in Appendix 1. These responsibilities may be fulfilled by an internal team or external contractor. Because many of the skills needed to act as an MSP for CommCare are the same as the skills required to act as an MSP for any large scale transactional platform, organizations should consider economies of scale in having the MSP team support more than one digital platform.

All engineers should be onboarded and be familiar with the CommCare cloud deployment documentation prior to starting the system installation. Each member of the team should review the prerequisites to setting up CommCare in production and should complete a quick install on a monolith server following the instructions detailed in the documentation.

Detailed documentation and instructions are available online:

- CommCareHQ Deployment Documentation
- Hardware Resource Requirements for CommCare HQ
- If working with a Managed Service Provider (MSP) please feel free to refer to the template Request for Proposals in the Appendix below

Checkpoint: Do not proceed to the next phase until the following have been completed

- MSP Engineers are onboarded with the CommCare cloud deployment documentation prior to starting the system installation. Each member of the team has reviewed the prerequisites and has completed a quick install on a monolith server.

2.4 Create Secure Migration URL and DNS Planning

During the application migration process users must stop syncing data to the development environment and change over to the production environment. To facilitate the process with minimal user interruption, changing the application to synchronize to a custom URL enables the MSP to toggle requests to the development environment, then after the migration, to the new server. This change occurs at the DNS level without requiring a change on each device.

The MSP sets up a domain name to be used for the migration, enables the application toggle, and releases a new build. This application change should be pushed to all users prior to initiating any data migration steps. Users who do not update their applications will be locked out once the switch to the new instance takes place. Detailed documentation for implementing this mobile application change are available online and should be completed prior to the local system installation.

The MSP is required to provide a domain name that they own and control, and will be required to add a special, temporary DNS entry in order for Dimagi's SSL certificate to work properly.

Detailed documentation and instructions are available online:

- Switch mobile devices to a proxy URL

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

□ Admin-level access to the mobile application
□ Domain name to be used for the migration, pointing to the old environment

Checkpoint: Do not proceed to the next phase until the following have been completed

- A new version of the mobile application(s) have been deployed and released to all users including the Proxy URL.

Phase 3: System Installation and CommCare Environment Setup

3.1 CommCare Installation

The MSP should now install the CommCare HQ software suite using the documentation and tools provided by Dimagi. After completing the installation, the MSP should run all applicable tests to ensure the CommCare environment is working as expected.
Detailed documentation and instructions are available online:

- Single Server Production Installation
- Multi-Server Installation
- Troubleshooting first time setup
- Testing your new CommCare Environment

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

- □ Name and job function of person who will be installing CommCare on the server
- □ Name and job function of person managing CommCare post-deployment
- □ What version of Ubuntu is installed on the server / cluster?
- □ What domain name directs to the server?
- □ What is the Public IP address for the server?
- □ What email gateway will be used for the application?
- □ What SMS gateway will be used for the application?
- □ What kernel livepatching service will you be using (e.g. Ubuntu advantage)? For Ubuntu advantage, what is the console output of running: `canonical-livepatch status --verbose`
- □ If you are using a hypervisor, what hypervisor software are you using? What is the response time of the support subscription you have (e.g. 1 business day)?
- □ What are the update / maintenance schedules for the software and firmware of the components of the hosting environment? Are all elements currently up to date with the latest available versions?

Checkpoint: Do not proceed to the next phase until the following have been completed

- Applicable tests as defined in "Testing your new CommCare Environment" are passed
- The system performs in real conditions.

3.2 [Optional] Link Cloud-Based Test or Development Environments

Some programs may opt to maintain a development or test environment in Dimagi's cloud environment in which users may build or test applications (referred to as an "upstream domain") before pushing them to the on premise installation of CommCare (referred to as the "downstream domain"). This is achieved by establishing a link between the Dimagi-hosted SaaS CommCare instance and the on-premise CommCare instance. Note that this link may be disconnected anytime in the future, and maintaining a development environment on the Dimagi hosted SaaS platform requires a paid subscription which includes support.

The MSP will provide information about the upstream and downstream project spaces, including the URL of the downstream local server project space(s) which will contain copies of the mobile application.

In order to proceed with establishing the link between each domains, the following must be established:

- A CommCare administrative user must have access to both the upstream and downstream project spaces
- The CommCare administrative user must have the Multi-Environment Release Management permission added to their role in both the upstream and downstream domains

Once the prerequisites have been met, the following commands need to be executed:

1. Run the `link_to_upstream_domain` management command on the downstream environment in order to link the upstream environment to the downstream environment
2. Sync the downstream environment with the upstream environment by going to Project Settings -> Linked Project Spaces -> select Sync & Overwrite for all the content wished to be synced to this downstream environment
3. Create a new “empty” application in the downstream project
4. Run the `link_app_to_remote management` command in the downstream environment in order to link the new downstream application to the upstream application
5. Update the downstream application by clicking the Update button when selecting this application from the UI on the downstream environment.

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

- □ URL of each of the project spaces which will be used for development or testing is accessible by an HQ administrative user
- □ URL of the production project space is accessible by an HQ administrative user

Checkpoint: Do not proceed to the next phase until the following have been completed

- HQ Admin user verifies that CommCare Application configuration may be pushed from a cloud-based development / test environment to an on premise production environment

3.3 [Optional] One-Time Migration of CommCare Application

Often application development commences on the Dimagi-hosted CommCare SaaS platform (referred to as an “upstream domain”) in parallel with deploying an on-premise instance of CommCare (referred to as an “downstream domain”. The application code must be moved to the on-premise instance of CommCare so that it can be pushed to end user’s devices. This is achieved by doing a one-time migration of the application from the SaaS environment to the on-premise instance before severing the link.

The MSP provides information about the upstream and downstream project spaces, including the URL of the downstream local server project space(s) which will contain copies of the mobile application.

In order to proceed with establishing the link between each domains, the following must be established:

- An CommCare administrative user must have access to both the upstream and downstream project spaces
- The CommCare administrative user must have the Multi-Environment Release Management permission added to their role in both the upstream and downstream domains
- Each downstream project space must have the "Linked Project Space" project setting enabled

Once the prerequisites have been met the MSP must:

1. Run the `link_to_upstream_domain` management command on the downstream environment in order to link the upstream environment to the downstream environment
2. Sync the downstream environment with the upstream environment by going to Project Settings -> Linked Project Spaces -> select Sync & Overwrite for all the content wished to be copied to this downstream environment
3. Create a new "empty" application in the downstream project
4. Run the `link_app_to_remote` management command in the downstream environment in order to link the new downstream application to the upstream application
5. Update the downstream application by clicking the Update button when selecting this application from the UI on the downstream environment.
6. Run the `unlink_apps` management command to cut the link between the upstream and downstream environments

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

☐ URL of each of the project spaces which will be used for development or testing is accessible by an HQ administrative user

☐ URL of the production project space is accessible by an HQ administrative user

Checkpoint: Do not proceed to the next phase until the following have been completed

- HQ Admin user verifies that the CommCare Application has been moved from the current to target project spaces
- HQ Admin user verifies that the connection between the current and target project spaces has been severed

Phase 4: Post-Setup Activities

Once you have completed the initial work to transition applications, it is also important to develop a plan to oversee the required long-term maintenance of the application software and infrastructure.

4.1 System Load Testing

Performance of the hardware infrastructure should be assessed to ensure that the system will perform under both normal and peak loads. The MSP should run performance benchmarking with the number of anticipated users under normal and peak loads. The system should perform under both normal and peak loads prior to going live. After baseline usage and load data is available, the MSP should refer to the commcare-resource-model to understand how to tune overall system performance and to anticipate potential issues related to scaling the system.

Detailed documentation and instructions are available online:

- Performance benchmarking for CommCare HQ using Locust
- Using commcare-resource-model

Checkpoint: Do not proceed to the next phase until the following have been completed

- Applicable load and performance tests are run and the system is deemed performant. Sample output for performance tests using Locust may be found here.

4.2 Configure Monitoring & Error Logging

Real-time monitoring of system activities is essential to understand how the system is performing at a given moment, to identify risks before the instance fails, and to troubleshoot issues as they arise. Monitoring helps to forecast resource requirements for future scaling, and alerts may be set up on various monitoring metrics to detect resource limits, anomalies which may cause server issues, or detect security incidents.

The MSP should install and configure monitoring tools for both error logging and monitoring CommCare's hosts and application / service indicators. The same tool may be used for both activities, or different tools may be selected. Sample dashboards and integrations are provided for DataDog, Prometheus and Sentry, but the MSP may select alternative tools.

Detailed documentation and instructions are available online:

- Monitoring and Alerting Infrastructure Metrics
- Setting up Sentry for error logging

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

- ☐ Link to system health monitoring tool / dashboard
- ☐ Link to error logging tool / dashboards

Checkpoint: Do not proceed to the next phase until the following have been completed

- System health dashboards are configured. Either pre-set dashboards are imported if using DataDog or system health dashboards have been created if using an alternative tool. Dashboards include the indicators outlined on CommCare Infrastructure Metrics.
- Error logging has been configured using Sentry or alternative tools using the documentation provided.
- Standard Operating Procedure (SOP) for monitoring infrastructure performance is established.

4.3 Test Backup and Restore Process

Dimagi recommends an offsite, ideally cloud-based, backup that is run automatically. Performing periodic system backups provides a means to restore the integrity of the system in the event of a hardware or software failure. Performing a full system backup and restore is a critical capability of the MSP and as such should be performed prior to it being relevant.

Detailed documentation and instructions are available online:

- Backup and restore process

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

- ☐ Link to the Standard Operating Procedure for conducting regular full-system backups including frequency of backups and testing plan
- ☐ What system will be responsible for creating backups for system recovery? What types of backups will exist?

Checkpoint: Do not proceed to the next phase until the following have been completed

- A full backup and restore is completed for the entire system including PostgreSQL, CouchDB, BlobDB, and ElasticSearch.
- A Standard Operating Procedure (SOP) is established for regular full-system backups.

Phase 5: [Optional] One-Time Data Migration of Collected Data

5.1 Data Migration of Collected Data

Note: This step requires support from Dimagi

A data migration needs to occur if it is expected that mobile workers will continue working without re-installing their applications, and if existing data is desired on the locally hosted system. In general the process for data migration requires strong collaboration between both Dimagi and the MSP, as the timing of the steps is important, and can only happen once the on premise system is up and running.

The MSP is required to provide a secure place for Dimagi to dump data to (e.g. a secured server which can be accessed by SSH), with enough disk space for a full data dump, and enough bandwidth to allow a timely transfer from Dimagi servers.

There will be a period of time between when the data dump is being created and when the data dump is imported to the new server when mobile workers will not be able to access the system (e.g. syncing, form submissions), so this expectation should be set with all stakeholders, including mobile workers. It is also recommended that a dry-run of the process is carried out to prevent any data loss and to iron out hardware or configuration issues prior to stopping the production server.

Importing a data dump is the MSP's responsibility. Detailed documentation and instructions to do this are available online:

- Transfer a Project From a Multi-tenant to Standalone Environment

In general, Dimagi suggests a two-part process to be followed when migrating collected data in order to test the process and minimize downtime for end users:

Phase 1: Trial Run

Dimagi provides a test file for a trial import. The MSP imports the test file, resolves any questions or issues, and then cleans the database.

Phase 2: Final Run

The MSP coordinates a window during which to block access to end users to prevent new forms from being synchronized. Dimagi generates a full export and pushes it to a secure location provided by the MSP. The MSP imports the full dump and reopens end user access to synchronize data collected while the system was offline.
A new, valid SSL certificate for the domain name will need to be provided and configured at this point.

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

- ☐ A secure location for Dimagi to dump the data export has been provided
- ☐ Name, email address and job function of person migrating data from current environment to production environment has been shared with Dimagi

Checkpoint: Do not proceed to the next phase until the following have been completed

- System downtime has been communicated to end users
- Request has been made to Dimagi to generate a data export
- Data has been imported to the new environment using provided documentation
- HQ Admin user verifies that the data migration has been completed successfully

Phase 6: Production Maintenance and Support

6.1 Ongoing Maintenance and Support

If you have reached this step, it means that you have successfully transitioned your cloud-hosted CommCare application to a self-hosted and administered server. Congratulations!

An organization endeavoring to manage its own CommCare server environment or working with a third-party to do so should plan for ongoing effort and system administration capacity not only in the initial phases of provisioning, setup, and installation, but towards the long-term maintenance of the application software and infrastructure as well.

At this point, common plans should also be in place that would apply to any digital solution, including disaster recovery and helpdesk escalation for infrastructure issues from end-users.

Prerequisites / Guiding Questions to Complete this Checkpoint

In order to complete all Checkpoints of phase, you will likely have to answer the following questions.

- □ Link to the Service Level Agreement for ongoing infrastructure support
- □ Name and job function of person who will be monitoring the forum and CommCare changelog
- □ Expected turnaround for both average and urgent changelog notifications
- □ Name and job function of person doing CommCare Cloud deployments
- □ What is the schedule on which CommCare will be updated and deployed?
- □ Link to Disaster Recovery Plan
- □ Link to issue triage process for receiving issues from end users

Checkpoint: Do not proceed to the next phase until the following have been completed

- Have established the appropriate Service-Level Agreement (SLA) for ongoing support
- Have professionals available to provide the required levels of support as outlined in the SLA.

Appendix 1: Example Request for Proposals (RFP)

Request for Proposal (RFP) to Support Transition to On-Premise Hosted CommCare Application

Project Summary

The open source digital platform CommCare (www.commcarehq.org) enables frontline organizations to build their own custom, permanent applications on a professionally managed foundation - at any scale, anywhere. CommCare was originally developed by Dimagi, and today has evolved into being one of the most widely deployed open source digital platforms and Global Goods in existence.

XXX Client XXX is currently running a Software-as-a-Service instance of CommCare as part of their digital health system. As this project scales, the XXX Client XXX is interested in transitioning their cloud-hosted CommCare application to a self-hosted and administered server.

The XXXClientXXX is seeking a Managed Service Provider (either internal or a third-party vendor) to support this process. The purpose of this RFP is to solicit support from an implementing organization to lead the design of this hosting environment for CommCare, and support the migration of the cloud environment to a third-party hosting environment.

Proposed Work

Current Environment

CommCare is an Open Source web application built in Python and Django. You can read more about CommCare on the Open Source page. XXX SaaS language XXX.
CommCare deployments typically use the global CommCare Cloud hosted instances of CommCare HQ. This approach removes the need for the project to manage its own CommCare instance and also guarantees service availability and access to the latest platform features and security compliance with HIPAA and SOC-2.
CommCare HQ is a complex, distributed software application, made up of dozens of processes and several pieces of third-party open source database software. It is designed to be run in a cloud data center consisting of servers running both the core web application code, alongside different Open Source or free services hosted on multiple virtualized servers. It has been built for performance, reliability, security, and scale rather than simplicity.

Proposed Environment

XXXClientXXX is seeking a Managed Service Provider (MSP) to provide options for migrating to a 3rd-party hosting environment. Bidders should review this document, "A Guide to Transition Your CommCare Infrastructure," which provides an overview of what this process should look like. This Guide links to all relevant resources and information, including Dimagi's documentation Commcare Cloud Deployment Documentation and the CommCare Cloud GitHub documentation.

A successful cloud environment architecture will specify the number, configuration, and size of the virtualized server hardware required, the configuration and requirements for local and external networking hardware, as well as the resources required for disaster recovery (DR) and the specified degree of service availability.

The design for the environment should also define Service Level Agreements which can be provided to bidders seeking to fulfill the roles required for successful long-term hosting. These SLA's should outline the expectations for what is needed by the bidder for the architecture's needs to be met.

Scope of Work

A successful Managed Service Provider shall be expected to fulfill the following Roles and Responsibilities.

No	Activity	MSP
1	Understanding Application Architecture	√
2	Design of Cloud Solution	√
3	Migration of application and other data from existing cloud to new cloud	√
4	Provisioning of support level or Equivalent for software licenses as mentioned in the RFP. Covering updates, upgrades, security patches, issue resolution at software level, bug fixing etc.	√
4	Configuration of Cloud Solution at proposed Data centers	√
5	Provisioning of the required hardware for Cloud	√
6	Network Connectivity between Primary and the DR site	√
7	Bandwidth provisioning at Primary and the DR site	√
8	Conduct periodic DR Drill	√
9	Cloud Service Provisioning through Self Service provisioning Portal	√
10	Operational and Functional testing	√
11	Functional acceptance signoff	

12	24x7x365 Support, Cloud service Provisioning, de- provisioning, updates, auto-scaling, security, firewall, anti-virus, bandwidth etc.	√
13	Maintenance & Management of Cloud Solution & infrastructure post implementation	√
14	Compliance to SLA’s	√

Local Environment Overview (If Open to Options)

The MSP will work with XXX Client XXX to determine options and provide cost estimates for the specified design. The MSP should source at N (ideally 3) options of hosting providers with cost estimates for the proposed design in terms of monthly running cost. Hosting providers must incorporate at least the following elements, with additional criteria as defined by XXX Client XXX to meet the needs of the environment:

Name	Description	Software
Firewall	The firewall prevents unauthorised access to the servers.	
Proxy / Load balancer	All requests to the system pass through the load balancer which handles the network connection and load balances the request to one of the web processors.	Nginx
Web request processor	This is the python web server that receives incoming requests and processes them. These would mostly be requests from the online platform such as user management, application configuration, reporting etc.	Python web service (e.g. Django)
Mobile request processor	Requests that come from the mobile phones are handled by this service. It receives incoming data and is also responsible for providing updated data to the phones.	Python web service (e.g. Django)
Web entry service	In some cases it may be necessary to do data entry via the web. This service handles those requests and processes them in a similar manner to how the mobile phone would.	Java service
Background task processor	There are many different background tasks that get run on this system both for data transformation as well as asynchronous processing of user requests e.g. data downloads.	Python task processor (celery)
Queue service	This services provides queuing functionality for the background task processor.	RabbitMQ

Primary database	This database houses the primary data set for the system. It is the authoritative source of data for the mobile phones and the web request processors.	PostgreSQL + CouchDB
File database / cold storage	The raw data received from the phones is stored here as XML/JSON documents as well as images, log files and possibly some data backups.	Binary object storage e.g. RiakCS / S3
Analytics database	This is a secondary data store for use in reporting and analytics. It contains a complete copy of the primary data in addition to a significant portion of the data in cold storage.	ElasticSearch, CouchDB
Cache	To improve performance the web and mobile processors make use of a cache service to avoid having to re-query the databases.	Redis
Change Log	A persistent log of data changes in the system. This is used for asynchronous processing and for keeping the secondary databases up to date.	Kafka

Local Environment Design

The MSP shall estimate the hardware required to run CommCare with the following scale estimates:

- [Number of Mobile Users]
- [Expected volume of Case Data per User]
- [Expected volume of Form Submissions per user per Day]
- [Other large / unique factors like Large Images to be submitted regularly]

Local Environment Cost Trade-Offs

The MSP will work with XXX Client XXX to determine the cost-benefit trade-offs across:

- Performance
- Up-time
- Support
- Availability, Redundancy, and Recovery
- Network, System, and Data Security

Selection Criteria & Instructions

Requirements

Any MSP must be able to support a well planned process to execute a successful transition and ongoing support of CommCare. In order to complete this Statement of Work, Dimagi asks that organizations have the following resources available:

- Technical Requirements: Have demonstrated experience meeting all of the criteria required to effectively transition CommCare's, as described in this Transitioning Infrastructure Environment document.
- Personnel Requirements: Managing a deployment of CommCare requires a team of 1-3 full-time DevOps engineers, depending on the size of the deployment. Key responsibilities are outlined below; detailed job descriptions and required skills are outlined in the deployment documentation here.
    - Cloud Application Management / Operations Engineer: Responsible for deploying the CommCare web application and dependent services, and keeping them up-to-date. Key duties include sizing and provisioning servers, debugging and maintaining individual services that are part of the cloud, supporting data migrations and monitoring the status and availability of individual services.
    - CommCare Application Administer: Responsible for configuring CommCare HQ from inside of the web application. Provides technical support for end users of the application and internal maintenance activities to be run from within the HQ web application.
- Experience Managing Open Source Tool: The organization has demonstrated experience managing an Open Source software platform. This includes active participation in the Community of Practice, contribution of code, and / or deploying or forking distributions of the codebase.
- Demonstrated Experience with Global Goods: The organization can point to experience deploying and maintaining other Global Goods or open source digital systems (DHIS2, OpenMRS or similar), and is willing to provide 3 external references about the process.

Period of Performance

This Statement of Work is designed to be completed within XXX time.

Submission Instructions

Submissions

A winning bidder will be expected to have the following deliverables:

- A matrix of each of the checkpoints in the Transition to On Premise Hosted CommCare document, with confirmation and examples of how you would go about completing this.
- Service-Level Agreements (SLAs) and cost estimates for Managed Service Providers involved in setup and run-and-maintain for the environment
- Cost estimates from N hosting providers for implementing the proposed design.

Submission Instructions

All materials should be submitted via email to XXX Client Email XXX , no later than XXX DateXXX

Points of Contact

- Name:
- Title:
- Email:
- Phone Number:

- Name:
- Title:
- Email:
- Phone Number:

Appendix 2: Skills Required to Maintain CommCare On Premise

In general, Dimagi considers the following list of items as to be under the purview of devops engineers working on CommCare:

- Deploys
    - Update CommCare code and dependencies
    - Database migrations
    - Done with Fabric (via commcare-cloud)
- Service provisioning configuration and management
    - Infrastructure provisioning
    - Service Installation
    - Service orchestration
    - Monit / Supervisord
    - Configuration
    - Upgrades
    - All done with Ansible (via commcare-cloud)
- Monitoring/Logging
    - Metrics and alerting
        - Datadog
    - Error reporting
        - Sentry (cloud / on-prem)
    - Logging
- Firefighting
    - System for on-call/fast response
        - Alerts
        - On-call rotations
        - Escalation
    - Automation or remediation
    - Retrospective
        - Root cause

- Futureproofing
- Documentation
- Knowledge sharing

- Performance
    - Monitoring
        - Metrics / APM
        - Collection of database usage patterns
    - Testing
    - Optimizing physical infrastructure
    - Optimizing services
    - Optimizing code
- Security
    - Best practices
        - Access, scope, configuration, etc.
    - Service updates
    - Library updates
    - Audit: holistically revisiting security practices
        - password rotation, audit log management, sudo access
        - audit report generation
- Scaling and Server Sizing
    - Horizontal and vertical
    - Sizing of infrastructure resources as the program scales
        - Forecasting
    - Optimization of Infrastructure services
- Backups
- Load testing
- Server sizing