

GFT > GET STARTED



Anthos



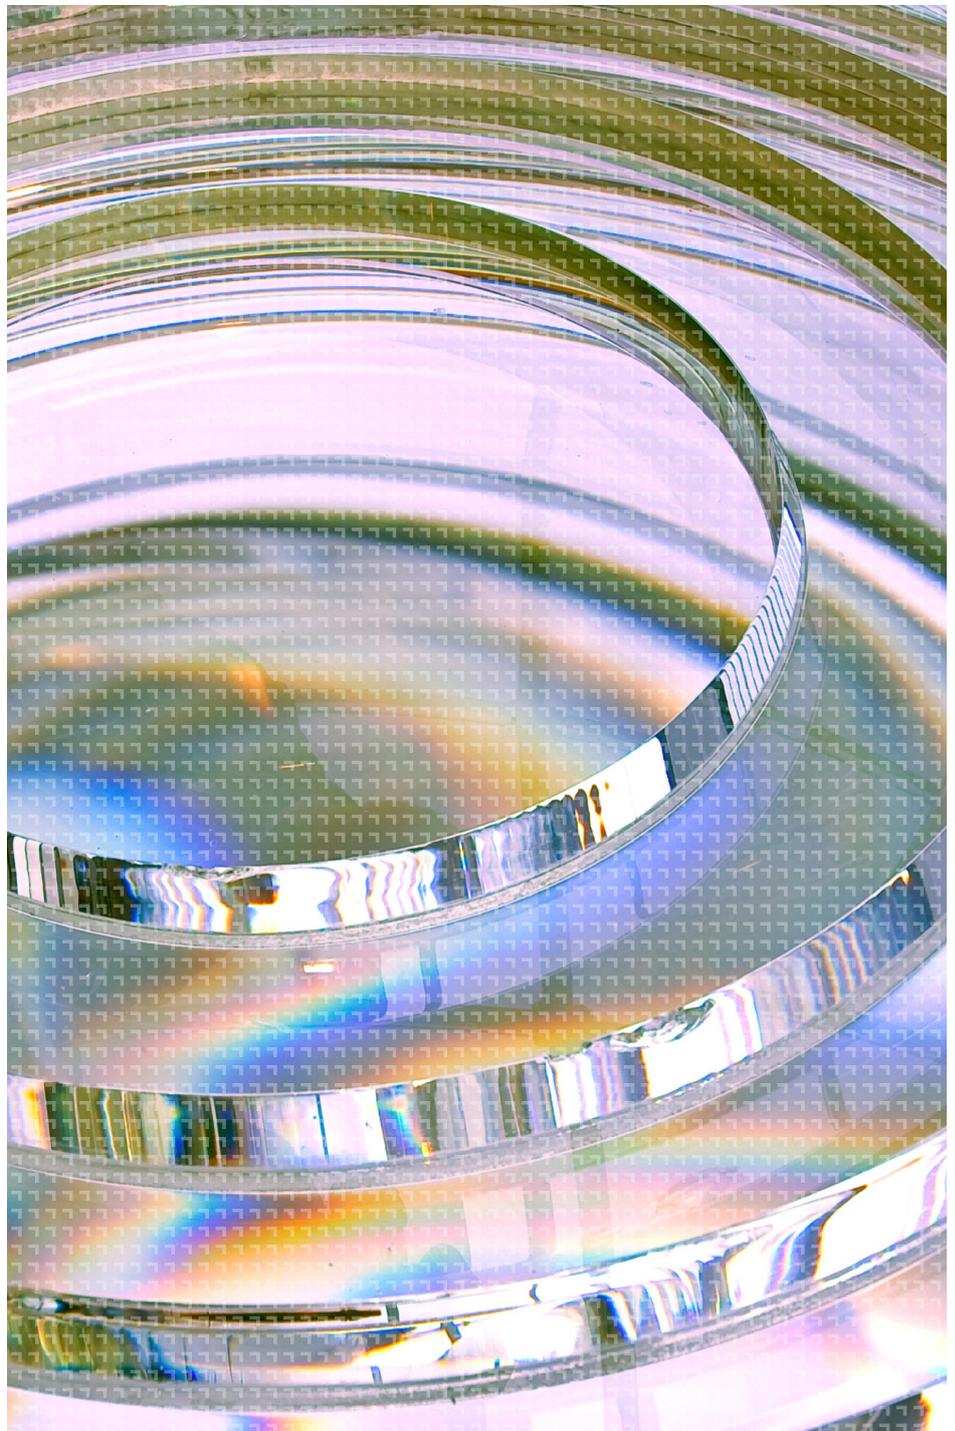
BREAKTHROUGH
**Partner
of the Year**
Europe, the Middle East,
and Africa
2019

Is this you?



Does your organisation have a technology environment that consists of both an on-premise and public cloud architecture? Are you a business or IT professional looking for a single management layer which will help make life easier, enable your organisation to be more agile, reduce total cost of ownership and achieve a much improved ROI? But struggling to choose the right path... you are not alone.

Cloud migration and the related digital transformation programmes can be both complex and time-consuming and many projects fail to even get off the ground. According to a recent analyst report, *'over half of IT transformation initiatives stall or are abandoned due to complexity, unfit processes or lack of skills.'* (IDG 2018). But do not let this put you off. In these volatile economic times, executing a comprehensive, robust cloud or digital transformation strategy is now more important than ever. This is where we come in because we can help to analyse your precise requirements to ensure success right from the start. Read on to find out how to get started.



The Anthos advantage: enterprise computing is multi-vendor, hybrid and multicloud



In a technologically fast-paced, dynamic world, with constantly changing requirements around performance, the customer experience (CX) geo-location and regulation, the lives of both business and IT leaders are becoming ever more complex and increasingly difficult to manage.

Google Anthos, delivered by GFT

This offers all organisations a unique ability to maximise much needed ROI, whilst enabling application modernisation and the rapid delivery of new capabilities at cloud scale, via a single management layer. Anthos is already helping to change the way enterprises view the cloud, delivering the ability to run applications with consistent development and operations experiences across multiple public clouds and on-premise environments.

Flexible pricing tailored to suit your specific needs

Anthos is a subscription-based service which is calculated on the number of 'virtual CPUs' your organisation needs, regardless of location (on-premise or in AWS, GCP or Azure). If there is a sudden spike in activity, its highly attractive 'elastic commercial model' ensures the revised burst pricing is charged per vCPU by the hour. This means users are only ever charged based on actual consumption.

GFT's Anthos Accelerator



Sound's great, right? But this is not the end of the story. Based on many years of hands-on experience, GFT recognises the requirements of each organisation's IT estate are unique. Hardware may be similar, with the usual cloud service providers already also be in place. However, the complexities of configuration,

compliance, security, networking, apps, DevOps pipelines and processes will all be subtly different. So, knowing how and where to start is never easy; both operationally and from a business case perspective. That's why we have introduced the GFT Anthos Accelerator.

With GFT's Anthos Accelerator, your company can quickly and securely establish a unified, production-grade presence across multiple cloud platforms.

This will pave the way for true mobility and cloud-scale for your modernised and future applications across all the cloud service providers (CSPs).

Experience the power of Anthos for yourself. GFT's Anthos Accelerator provides a secure, production-ready landing zone on multiple clouds, all delivered within a single tailored work package.

Flexible engagement model



SMALL

The GFT Anthos Accelerator Lab is entirely hosted by GFT in our public cloud and VMware environments. Here, you will meet with expert cloud engineers and will meet with our highly experienced team of technical consultants and cloud engineers who will help you to understand precisely what you need to do to deliver a successful outcome for your business.

MEDIUM

This particular GFT Anthos Accelerator Lab session is pre-provisioned in a Google Cloud project owned by the client, with one or more additional available clusters added from other client-owned clouds. Benefit from a highly interactive meeting, where you will learn from GFT experts how to create a successful hybrid, multicloud environment in an Anthos Lab that is tailored to your precise business needs. We are here to help.

LARGE

This is a unique opportunity to work with GFT specialists in a suitable, production-ready VMware environment. Learn all you need to know about how Anthos GKE OnPrem can be installed and integrated with a client-owned GCP project to provide a true hybrid environment, upon which the lab can be delivered. Get ready to do better business.

GFT's Anthos Accelerator delivers reference hybrid/multicloud Anthos implementations on the platforms of your choice – all within a single, flexible, guaranteed commercial package.

Options include:

- Production-grade, secured Google Kubernetes Engine clusters in multiple Anthos-supported environments
- Integration with IAM, Anthos Service Mesh and Anthos Configuration Management
- Operational dashboards for observability and monitoring
- Training, consultation and support as required.

Get started with Anthos



Before the GFT Anthos Accelerator can be deployed, the first step is for GFT to undertake a detailed, but objective evaluation of your organisation's needs and IT estate. This includes:

- A viability and value assessment of the initial applications you should focus on, and the reasons why
- An overview of what is needed to containerise these applications, if at all
- An analysis of your organisation's CPU count of your existing IT estate, on-premise and in the cloud
- An analysis of the average CPU performance/usage across your business.

Statement of work

Based on the results of the evaluation process, GFT will present a fair and accurate statement of work, including detailed costs for the deployment of the Anthos platform and the associated services.

Our Anthos credentials



- GFT is a Google Cloud Premier Partner and one of a few global technology service companies with real-life Anthos deployment experience
- Google Cloud EMEA Breakthrough Partner of the Year 2019
- 325+ Google Cloud certified engineers and technicians
- Accredited and experienced Global Anthos Service Partner
- Experts in delivering cloud projects within highly regulated environments.



Contact details



Ben Jones ■ Cloud Sales Manager
Mobile: +44 7890 386 367
Email: [Ben Jones](mailto:Ben.Jones@gft.com)

About GFT



GFT Technologies SE is a global technology partner focused on digital transformation. Founded in 1987, we have over 6,000 specialists in 15 countries.

[gft.com](https://www.gft.com)

GFT helps transform a global exchange with multicloud adoption using Anthos

Background

Our client, a global trading exchange based in London, had identified that its current technology infrastructure was a significant inhibitor to business efficiency and growth. Specifically, they wished to exit their existing data centre, which required all corporate web applications (Tier 0 with RTO and RPO of 0) to be containerised and moved to the cloud running in container orchestration platforms. This required orchestrating workloads across various different clouds (GCP and AWS) and on-premise systems utilising VMware. Despite there being a tight deadline for the project (a six month timeline), the goal was also to re-architecture the applications in order to leverage the maximum benefits from the containerisation process.

The goal was to migrate all applications from the legacy data centre with zero downtime over the course of the project, thereby optimising the IT portfolio, with the new platform being configured for maximum cost-efficiency, keeping on-premise capacity at maximum and orchestrating workloads to the cloud only when absolutely necessary.

RPO = Recovery Point Objective
RTO = Recovery Time Objective



Engagement overview

GFT established two connected teams (pods) responsible for:



GFT Pod 1 – container orchestration and re-platform buildout

GFT Pod 2 – application containerisation and migration

The teams worked closely with the client to define the security posture for containers as well as building the fully automated Anthos GKE platform and ‘migration factory’ for all corporate web applications, with the goal of all applications being migrated off the legacy data centre with zero downtime.

Overall engagement objectives:

- Data centre exit
- Buildout of new container orchestration

platform capable of orchestrating workloads across multiple CSPs (AWS and GCP) and on-premise VMware environment

- Re-platforming and containerisation of all (30+) corporate web applications
- Defining security posture for containers across the whole organisation
- Zero-downtime migration of multiple Tier 0 applications (with hard requirements RTO and RPO requirements of 0)
- Platform cost optimisation – maximise on-premises capacity and utilise cloud-based pods only when absolutely necessary based on a config-driven rules engine.

Added value objectives



Reduced blast-radius

By adopting Anthos, the client would be able to significantly reduce the ‘blast-radius’ of their corporate web application suite, making any failure of domains easy to understand and analyse, by treating both compute instances and entire clusters as ‘cattle’ which can be swapped out as required. The objective was to separate workloads by hosting them on physically and geographically independent Kubernetes Anthos GKE clusters managed from the same interface. Furthermore, having a consistent application development platform and automated policy enforcement helps the client developers’ code to be deployed with much greater confidence – significantly reducing the risk of errors and reworks that have a customer impact.

Increased development velocity

Utilising the Anthos GKE platform significantly improves development productivity, with the client development teams being able to spend more time coding, testing and experimenting with new functionality rather than on deployment and configuration activities. With the adoption of a GitOps approach, the client’s platform team is able to more quickly configure, patch and update platform components, release software, and migrate applications into Anthos across the hybrid environment.

Enhanced security

By providing a consistent and unified security policy creation and governance interface, Anthos is helping the client’s security team to boost productivity by simplifying security policy creation, deployment and enforcement. Anthos has been configured with several features to help secure the client’s workloads, including the contents container image, the container runtime, the cluster network and access to the cluster API server. With Anthos Config Management, the client is able to define and enforce policies in clusters both on-premise and in the Cloud Service Provider (CSP) of their choice. Application security has been delivered by design, significantly reducing risk of vulnerabilities and helping address them more quickly when discovered.

Mitigated risk of vendor lock-in

By providing maximum application portability across the different public clouds and any VMware-certified on-premise infrastructure without specialised hardware, it is possible to reduce dependency on specific cloud providers or even infrastructure vendors. Since the Anthos platform was built using market-leading open source components, the client can architect applications in Anthos with a lower barrier to adoption of future technology based on those components, without fear of vendor lock-in.

While the current Anthos GKE platform is dependent on VMware, more platforms are on the roadmap, with bare metal support (a key request from clients), scheduled for Q3 2020 – significantly reducing the relatively high costs associated with VMware licensing and further diversifying the client’s infrastructure portfolio, reducing risks of costly vendor lock-in.

Reduced costs

By adopting unified container orchestration tools built on open-source technologies, the client can significantly reduce tooling costs and operations labour that was otherwise required to manage multiple separate development environments, tools and platforms. Sourcing, integrating and deploying multiple open-source tools and services in an ecosystem as vast as Kubernetes had proven very costly for the client in the past. The Anthos platform has packaged and delivered these technologies, fully automated and ready for use, across all the required environments and to all of the development teams.

Additionally, in the process of application modernisation and containerisation of workloads, the re-architected applications have become significantly more cost-efficient to run and operate, by using the underlying hardware and cloud infrastructure more efficiently, by optimising container performance and resource utilisation.

The platform has been configured to maximise the remaining on-premise utilisation and utilise applications in the cloud only when absolutely necessary (e.g. due to demand spikes or hardware failures), further reducing overall running costs.

Delivery overview

GFT Pod 1 – container orchestration and re-platform buildout



Whilst GFT deployed two discrete teams (Pod 1 and Pod 2) who worked independently on the project, there was a great deal of collaboration between the pods as the project progressed. The detailed project stages and deliverables are outlined separately below. However, both pods progressed their distinct tasks simultaneously, in order to create maximum efficiency and optimal time-to-completion, given the aggressive timescales for the project.

Sprint Zero (0): Onboarding, environments discovery, licencing, tooling setup

Sprint Zero activities included on-boarding the team into the client organisation, completing all necessary training, all relevant tooling setup (e.g JIRA, AWS workspaces, GCP console, Source repositories), working with procurement on approval for the trial Anthos licence, as well as mapping out all the key client stakeholders to work alongside and the requirements to successfully deliver the platform.

Sprint 1: Workshops, definition of solution architecture, GCP pre-requisites, on-premise requirements

In Sprint 1, the team focused on organising and running various workshops, including: containers, Anthos GKE, security, network topology as well as DevSecOps development and deployment processes. Also essential was the creation of a detailed solution architecture document outlining the platform which required approval by various client stakeholder bodies – architecture board, compliance board, security board. Additionally, the team created all the necessary GCP projects and put in place the required guardrails in order to be able to start testing Anthos deployments. The team also inspected the on-premise setup to understand all of the necessary changes required for the installation of the GKE on-premise solution.

Sprint 2: Container registry setup, DevSecOps pipelines design, base image requirements and buildout

As the pilot application has a RTO and RPO of 0, in Sprint 2, the team needed to provision and set up an adequate highly available container registry to meet the above requirements. The Elastic Container Registry (ECR) was selected and provisioned for early testing and the DevSecOps pipeline for building and deploying containers was designed and approved by the client architecture board. Base image requirements were defined and an early build pipeline was provisioned.

Sprint 3: AWS VPC creation, network and pod security policies definition, Terraform deployment validation, continuous integration for Anthos management cluster, manual user cluster deployment

In Sprint 3, the team set up all necessary AWS VPCs as per the client requirements, worked with InfoSec on advising and defining all necessary network and pod security policies required for production deployment, defined a Terraform validation framework and built an automated continuous integration pipeline for the deployment of a GKE management cluster in the client AWS development environment. The team also provisioned the very first user cluster in the development environment to enable initial connectivity and configuration testing.

Sprint 4: Jenkins slave setup, GCP security controls, OpenID Connect (OIDC) and Hashicorp Vault setup, management cluster destroy pipeline, user cluster build pipeline and provisioning for first development team, application load balancer setup

In Sprint 4, end-to-end connectivity was established for Jenkins slaves to provision all necessary clusters for the first development team. GCP security controls were set up, including service account creation and OIDC setup for secure connection to GKE dashboards in GCP. Hashicorp Vault was integrated into the GKE clusters for keys and secrets management. Automated management cluster destroy pipelines were provisioned for efficient clean-up of resources and the first user cluster was provisioned using an automated Jenkins pipeline in the AWS development environment. An application load balancer was manually provisioned to enable the development team to start their regression and functional testing.

Sprint 5: Production environment set-up, BlackDuck container scanning, Anthos Configuration Management implementation and deployment, CrowdStrike integration, Datadog daemonSet, Dome9 integration

In Sprint 5, the team worked with the client to provision the AWS production environment, updated the existing DevSecOps pipelines to include container scanning using BlackDuck, and rectified all the known Common Vulnerabilities and Exposures (CVEs). Anthos Configuration Management was installed and configured to monitor the network and pod security policies defined in Sprint 3. Automated steps integrating CrowdStrike, Datadog and Dome9 daemonSets were added to the DevSecOps pipeline.

Sprint 6: Management and user cluster production deployment, regression testing, scalability testing, penetration and information security vulnerability testing

In Sprint 6, the team adjusted the continuous integration continuous deployment (CI/CD) processes to point to new production environments and deploy both management and user clusters in a production-grade configuration. All the necessary unit, functional, regression and scalability testing were performed ahead of the scheduled go-live.

Sprint 7: Parallel run, operational runbooks, deployment cutover

In Sprint 7, the team focused on running the pilot application on production clusters in parallel to the existing live version, while preparing operational runbooks to handover to the client's operations team. Zero downtime deployment cutover had been performed, whilst the team monitored the state of the application using the GCP Anthos dashboard plus other monitoring tools.

Sprint 8: Handover to operations, roadmap curation

In Sprint 8, the team delivered a series of workshops and handover sessions with the client operations team to familiarise them with the provided operational runbooks and support in the early weeks of platform maintenance and operational work. In parallel, the team met with the client to define a roadmap for migration of the remaining applications, potential platform improvements and additional requirements requested by the client throughout the project.

GFT Pod 2 – application containerisation and migration



Sprint Zero (0): Onboarding, tooling setup, workshops

In Sprint 0, the team proceeded with onboarding, completing all necessary training, requesting access to services and tools such as Bitbucket, JIRA and Confluence. The team also organised a series of workshops with key client stakeholders to familiarise them with the data centre exit roadmap strategy, containerisation process, shared security model, and the benefits of containers and container orchestration platforms such as Anthos GKE.

Sprint 1: Containerisation skeleton, golden image

In Sprint 1, the team began familiarising itself with the applications requiring containerisation and migration, and built a skeleton of a similar containerised application in a specially created Anthos Lab environment. This skeleton included a Drupal pod with php-fm and nginx deployed to match the specification of the client application.

Additionally, the team worked with the client central engineering team to understand the base image provided and created a Continuous Integration (CI) pipeline to create a 'golden image' i.e. an enriched image with all the required dependencies for a group of web apps being migrated.

Sprint 2: Solution architecture document, app image build pipeline, health-checks, quotas, namespaces, secrets integration, Kube-scan integration

In Sprint 2, the team worked with the client to create a detailed Solution Architecture Document (SAD) that was then approved by the various client boards (architecture, compliance, security) in order to gain access to the client AWS development

environment. The team enhanced the application skeleton built previously by including necessary health-checks, quotas, namespace creation and secrets integration with Vault. Kube-scan was also integrated to enable Kubernetes vulnerability scanning.

Sprint 3: App containerisation, app deployment on Amazon Elastic Kubernetes Service (EKS), Application Load Balancer (ALB) integration, nginx golden image pipeline

In Sprint 3, the team successfully containerised the pilot application and deployed onto an EKS cluster in order to start functional testing of the application whilst waiting for the Anthos GKE user cluster to be created by the platform team. The EKS deployment was configured with an application load balancer to allow ingress from the public internet. A continuous integration pipeline was created to automatically build a nginx golden image, and nginx was deployed as a side-car to allow internal traffic.

Sprint 4: Datacentre exit roadmap, Terraform deployment, helm charts

In Sprint 4, the team defined a roadmap and a 'migration factory' approach for the remaining web applications in order to enable the 2020 data centre exit objective to be achieved.

They also worked on finalising the terraformed deployment of auxiliary AWS resources (persistent storage, ingress, firewalls) as well automating application deployments using helm charts.

Sprint 5: Pod security and network policies implementation, application regression testing, SMTP integration, Common Vulnerabilities and Exposures (CVEs) rectification

In Sprint 5, the team focused on implementing both pod security and network policies according to the clients information security requirements, completed the final SMTP integration and the pilot application was deployed to the non-production environments using automated CI/CD Jenkins pipelines. A fully automated regression testing suite was implemented to ensure functionality was as expected and all the CVEs were rectified.

Sprint 6: Quality assurance, penetration and scalability testing, production environment deployment, cutover plan

In Sprint 6, the pilot application was deployed to the client's production

environment and quality assurance, penetration and scalability testing were performed. A detailed production documentation was produced and aligned with customer internal processes, and a cutover plan was finalised and agreed on.

Sprint 7: Parallel run and and cutover

In Sprint 7, the legacy application was run in parallel with the newly containerised application running on top of Anthos GKE. A small number of clients with special alpha access tested it and confirmed all necessary functionality was behaving as expected. At the time of cutover, DNS propagation was established and the application was switched to the new version with zero downtime.

Sprint 8: Monitoring and, migration of remaining applications

In Sprint 8, the team closely monitored the logs of the application to ensure any issues were managed swiftly. However, no issues were have been reported in the week after go-live which is a testament to the thorough planning and testing approach taken. All necessary handovers were given to the client operations team, and after celebrating the go-live, the GFT pod continued the engagement to containerise and migrate the remaining client applications.

Outcome for the client



Client challenge, goals and objectives	Delivery outcome
Buildout of new container platform capable of orchestrating workloads across multiple CSPs	Created a robust infrastructure utilising Anthos, capable of managing AWS, GCP and an on-premise VMware environment
Re-platforming and containerisation of all (30+) corporate web applications	Created a migration factory approach for systematic migration of all core web applications in a structured and organised manner
Define security posture for containers across the whole organisation	Ensured the highest quality of control and security for applications in the new multicloud environment applying industry-leading best practice
Zero-downtime migration of multiple Tier 0 applications	Achieved the establishment of the environment and migration of applications with zero-downtime
Platform cost optimisation	Improved the efficiency of IT infrastructure by moving to a hybrid and multicloud environment and by reengineering applications to operate more efficiently
Data centre exit	Enabled the client to establish a new multicloud environment that will ensure the data centre exit strategy will be achieved as planned
Define a network operations role	These users in Tranquility Base can only view the Network Shared Services.
Define a systems operations role	These users in Tranquility Base can only view the Operations Shared Services.



About GFT



-  blog.gft.com
-  twitter.com/gft_en
-  linkedin.com/company/gft-group
-  facebook.com/GFTGroup
-  gft.com

GFT Technologies SE is a global technology partner focused on digital transformation. Founded in 1987, we have over 6,000 specialists in 15 countries.

gft.com