

**Hybrid Cloud: The New Design Center
Questions and Answers from the Webinar**

GDT WEBINAR SERIES

HYBRID CLOUD IS THE NEW DESIGN CENTER



Webinar Questions & Answers

Thank you for attending our webinar on June 14th, *Hybrid Cloud is the New Design Center* in partnership with Cisco and NetApp. We have collected the unique or representative questions from the chat and have answers to all below.

Please do feel free to follow up with us on this topic or any other or if you would like to go deeper into any question/answer that follow.

How many people can access FlexPod at a time?

There are a number of ways to answer this question. In terms of infrastructure scale, FlexPod supports up to 20 UCS chassis and 160 UCS servers in a single UCS management domain. In terms of an example of tested/validated configurations, FlexPod has been tested to support ([CVD here](#)):

- 6,700 users based on the Windows Server 2019 sessions
- 5,400 instant clone non-persistent virtual Windows 10 desktop users
- 5,400 full clones persistent virtual Windows 10 desktop users
- 5,800 mixed users

This example is nowhere near the limit of a single domain from a management perspective as it is only using 32 x UCS M5 generation servers. Using newer UCS-X servers (M7) with higher consolidation ratios you can see that the scale of FlexPod will support most any workload or set of workloads.

If you would like to review scaling to more than a single UCS Manager Domain, we would be happy to discuss how UCS Director and Intersight can be leveraged to support your needs.

How often do we need to replicate that data to the cloud?

Replication from on-premises to cloud has a number of factors to consider. Once a baseline copy has been established, replication timing or schedule will be determined by a number of criteria:

- Change rate of the base/source data
- Recovery Point and Recovery Time Objectives for the data in consideration
- Available bandwidth
- Latency of the connection
- Compressibility of the data in flight

Once we have an idea of the above factors, we would help identify the necessary solution elements and components needed to meet your needs. In some cases, once per day updates are sufficient and in some cases the updates are near real time.

Is Chat GPT reliable?

Chat GPT certainly has the ability to be a useful tool... Asking Chat GPT if it is reliable yielded the following response:

As an AI language model, I'm based on OpenAI's GPT-3.5 architecture, which is known for its impressive language processing capabilities. However, it's important to note that the reliability of any AI model, including GPT-3.5, depends on several factors.

1. Training data: GPT-3.5 has been trained on a vast amount of data from the internet, which helps it generate responses based on patterns and information present in that data. However, the model's responses are only as reliable as the quality and accuracy of the data it has been trained on.
2. Knowledge cutoff: My training data goes up until September 2021, so any information or events that have occurred after that date are unknown to me unless I have been explicitly updated. This limitation means that my responses may not reflect the most current information available.
3. Bias and errors: AI models can inadvertently exhibit biases present in the training data, leading to potentially biased or inaccurate responses. OpenAI has made efforts to mitigate bias during training, but it's impossible to eliminate it entirely.
4. Lack of real-world understanding: Although GPT-3.5 can generate coherent and contextually relevant responses, it does not possess true understanding or consciousness. It relies on statistical patterns in the training data rather than genuine comprehension.

To ensure reliability, it's essential to verify information from credible sources and use critical thinking when interpreting AI-generated responses. AI models like GPT-3.5 can provide valuable insights and assistance, but human judgment and verification are still crucial for making informed decisions.

It is our opinion that Chat GPT is a useful tool but that the use of the output should be reviewed by subject matter experts as a matter of practice. For the purposes and use of the Chat GPT output in our presentation, it is, after our review, accurate and appropriate for the intended use in our discussion. 😊

How will our data be secured? Will that be in the next webinar?

There are a number of variables on securing data. Within the context of FlexPod, the full stack architecture does contain and support multiple security elements. Please see slide #26 in the presentation for a view of this. We also believe that security has elements of prevention, detection, and recovery and FlexPod is certainly an architecture that supports these elements natively and with additional 3rd party integrations. We would be happy to discuss your security requirements in more detail and in the context of your specific environment and security goals.

See the answer below for future topics for webinars.

Will this session be recorded for viewership at a later time?

Yes, we will provide a link to a PDF version of this presentation and an on-demand video.

Is it cheaper to get my CVO (Cloud Volumes ONTAP) licenses from GDT than buying on-demand from the cloud provider(s)?

There are free trial versions of CVO available to test and perform proof of concept/value with. We have found that procuring your licenses/subscriptions for CVO are more cost effective (e.g., cheaper) when secured through GDT.

How does the Cisco and NetApp partnership address data management and security concerns in a hybrid cloud environment?

On the data management aspect, NetApp provides a consistent, predictable, and secure set of data management services that work in the data center/colo, at the edge and in and across all three major cloud service providers (AWS, Azure, Google Cloud). NetApp supports encrypted data at rest, encryption of data in flight (replication), multi-factor authentication, multi-admin verify for destructive or risky operations. NetApp also provides to their customers AI driven recommendations in the customer accessible Active IQ environment and conforms to multiple security standards such as Common Criteria, NIST FIPS 140-2 and is included in the DoDIN APL approved product list. Unique to NetApp for a storage vendor is the NSA CSfC validation for hosting NSA top-secret data.

For Cisco, they have deep and wide portfolio of security products that cover network security, user and endpoint protection, intrusion prevention, email security, VPN, Web security, workload security, MFA, single sign on, advanced malware protection, firewalls, DDoS protection, router security and more.

Both NetApp and Cisco have AIOps driven security capabilities, API integrations with multiple 3rd party solutions and tools.

What active security measures does GDT implement to ensure security?

In terms of active security measures in the context of cyber security vs. physical, and we do have physical security controls in place as well, we follow generally acceptable security controls as documented by NIST, etc., to identify assets and risks, protect, detect, respond, and recover.

These capabilities are also represented in our selection of security focused technology and partner companies.

We have a team focused on cyber security that can to consultative, advisory services, workshops, table-top exercises, penetration testing, remediation, and managed services for security. We would love to discuss your security questions, needs and inquiries.

What topic will the next webinar most likely cover?

Topics for future webinars will include as it pertains to Hybrid Cloud

- Ransomware prevention, detection, and recovery
- Cloud Repatriation – why and how
- The “yottabyte” problem – how data growth is the hidden driver of cost for all of us
- VMware and Hybrid Cloud
- Full Stack Observability

Additional topics that are coming are on topics in the areas of

- Asset and lifecycle management
- Security
- Modern networking
- SD-WAN as a Service

We are always looking for relevant topics, please let us know what you would like to hear!!

How does this all-in-one solution compare in TCO to the piecemeal solution?

(Question was also accompanied by the comment: “TCO to me is purchase, maintenance, renewals, decommission, disposal.” We agree 😊)

Each environment will have different results. If refreshing older environments characterized by compute and storage elements that are 3 years or more older, TCO as an example compared to UCS M5 based systems compared to newer UCS-X/M7 based systems is estimated to be 23% lower for a typical VDI solution. In terms of network management, you can expect to see approximately 20% less time in management and maintenance tasks, \$3M average cost savings in hardware refresh and infrastructure costs and up to a 4x reduction in downtime.

How many software scripts can you support?

We do not see a practical limit to the number of scripts that can be supported. Having said that, we do believe that in most organizations there is benefit from a script matrix that is regularly evaluated and updated to avoid duplication of functionality where it may not be needed. Intersight does have a RESTful API that supports the OpenAPI Specification (OAS), supports Python and PowerShell via SDKs that enable integration with Ansible, Chef, Puppet and other DevOps and IT Operations Management (ITOM) tools. If you are a ServiceNow shop, you can ingest data from Intersight into the ServiceNow CMDB which allows visibility to your Intersight connected inventory in the ServiceNow IT operations management platform.

Could you provide the pricing and other benefits in depth?

Pricing is variable as each customer deployment is typically sized for their specific compute, memory, connectivity, storage, and performance requirements. Also consider that each customer has different strategies as it pertains to growth, headroom, and consumption factors, e.g., depreciation schedule, overhead, lease, as a service, etc. We would love to discuss your specific needs and find the right solution for you.

As to benefits in depth, we would offer the following for consideration, keeping in mind that an exhaustive deep dive would be a lengthy document.

- Scalability – The ability to scale resources effectively and efficiently for compute, storage and network as needed and to do so in a cost-effective manner.
- Availability – FlexPod offers redundant components, resilient architecture, and the ability to extend across multiple sites.
- Performance – FlexPod supports a mix of performance options. In terms of high performance, FlexPod supports the latest Intel, AMD and NVIDIA CPU/GPU processors, high bandwidth low latency networking, end-to-end NVMe storage.
- Versatility – FlexPod has greater than 200 Cisco Validated Designs and can within the same architecture support growth from a small, entry level configuration through to the highest demanding, highest performance deployment. The architecture does not change, management stays the same. No compromises.
- Simplicity – One management interface can handle deployment, configuration, Day 1 and following with APIs available to support any Infrastructure as Code strategies you may pursue.
- Security – Cisco and NetApp are highly focused on security at all layers of the stack. At the host level, network, and storage. FlexPod can seamlessly integrate into your existing security framework(s) and has multiple built in security features.
- Cloud – Both Cisco and NetApp have multiple cloud integration elements. In the case of FlexPod, NetApp can extend your data and data services to any of the main cloud service providers.

Is it Intersight that is leverage for UCS management? Does that extend to the architecture as a whole? Automation and Orchestration?

Yes, Intersight is the central point of UCS management. You still do have UCS manager as part of the solution. Intersight is used in the case of FlexPod to manage the entire storage, network and compute configuration with observability and monitoring. Automation and orchestration are part of the workflows within Intersight, and the solution supports common automation tooling such as Ansible, Terraform, etc.

Is there a storage limit?

Within a single NetApp cluster, the largest cluster supports up to ~115PB useable capacity. Multiple clusters can be addressed by a FlexPod. This does not factor in storage efficiencies such as deduplication, compression, etc. or the ability of NetApp to tier cold data. NetApp supports 300 TB volumes, 127 TB LUNs and up to 20 PB in a single file system.

What are some best practices or recommendations for organizations embarking on a hybrid cloud modernization journey with Cisco and NetApp?

Great question! We typically recommend a number of initial / first steps. If modernizing on-premises, we do want to go through a sizing and dependency mapping exercise to ensure that the new solution is aligned to the business strategy. Done correctly this allows us to recommend a service catalog based on workload requirements. This normalizes the discussion to cost per GB per month for each service level. This is done with an analysis of the actual workloads and then helps rationalize what the cost and benefit matrix is for how the architecture is built initially and answers the question of what the same workloads would cost in and across clouds. We would recommend a security review as well as a data classification exercise. Finally, we would review current backup and recovery in place and make recommendations.

Is it best to be in a “cloud adjacent” Co-Lo?

This is a question often encountered where putting data into the cloud is viewed as not desirable due to security, data governance (e.g., GDPR, etc.) or for performance reasons but where cloud services are of value. This is a valid and often best balance for customers. If the Co-Lo provider has for example AWS Direct Connect and Azure ExpressRoute connectivity, it is possible to you may be able to have one set of data accessing services in both clouds. There are differences in this example between the AWS and Azure offerings that may be relevant to a specific deployment scenario and should be evaluated. This would be an area where GDT would want our Cloud team to engage and validate.

How can GDT streamline your experience and provide a bespoke hybrid cloud solution that meets your specific needs>

GDT has the experience with Cisco and NetApp to properly architect your solution based on your environment, goals, and requirements. We do this day in and day out across our customer base.

As stated earlier, we prefer to start with a more consultative evaluation of the current state, desired outcomes, and alignment to your business strategy. Understanding your cloud experience as well as plans does factor in.

We typically see that the recommended solution for on-premises helps standardize the environment, making automation more feasible, operations simplified and a more reliable infrastructure. As it relates to hybrid cloud, in addition to optimizing and modernizing on-premises we must also look at any existing or planned cloud adoption and evaluate security, R-lane analysis and a FinOps approach and practice around cloud service consumption.

Can you expand on the benefit of leveraging this solution with an Application Performance Monitoring Tool? Also, is their better integration?

There are real benefits to extending observability and operations to include applications. If we consider that most organizations are really focused on their applications and data, then this is essential to provide a world-class digital experience.

In this example with Cisco Intersight, we would look at the combination of Cisco's Intersight Workload Optimizer with AppDynamics. AppDynamics gives you the deep observability into the application and increasingly the other applications that is being accessed by the one application. AppDynamics can optimize performance, improve security, and help in time to resolution for issues at the application level.

When adding in Intersight Workload Optimizer to the equation, you can now create an experience journey map which leverages machine learning to ingest data generated by your applications and AI and connect that to the underlying infrastructure. This approach allows you to quickly identify dependencies, collaborate across teams and either resolve issues faster or deliver applications faster.

Not sure that there are better integrations... but we are always open to learn!