

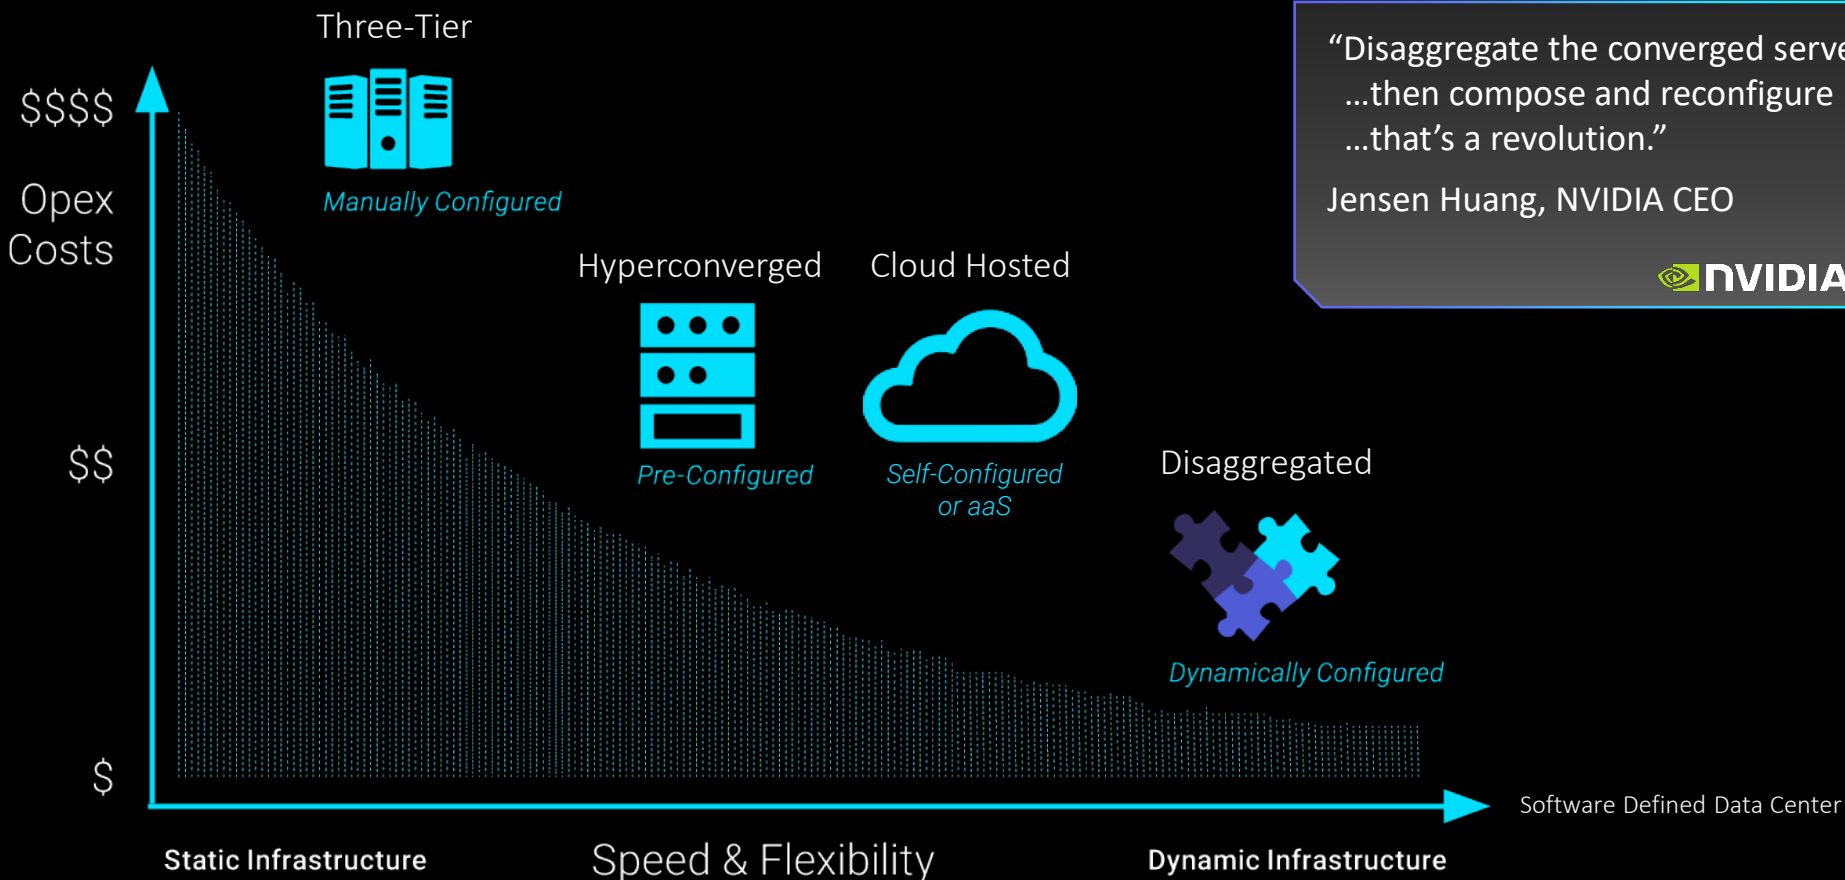
# Breaking out of the Physical: Composable infrastructure of the software defined server

Scott Houppermans

Senior Solutions Architect – Federal & Enterprise




# Datacenter Transformation




“Disaggregate the converged server  
 ...then compose and reconfigure  
 ...that’s a revolution.”  
 Jensen Huang, NVIDIA CEO



Source: Gartner



# Worldwide Cloud Spend & Utilization Rate



Source: <https://bit.ly/3uYLoUs>  
<https://tinyurl.com/4fy93zfq>





# Cloud Service Complexity

## COMPUTE

Cloud Functions  
App Engine  
Cloud Run  
Kubernetes Engine (GKE)  
Compute Engine  
Bare Metal Solution  
Preemptible VMs  
Shielded VMs  
Sole-tenant Nodes  
VMware Engine

Event-driven serverless functions  
Managed app platform  
Serverless for containerized applications  
Managed Kubernetes/containers  
VMs, GPUs, TPUs, Disks  
Serverless for containerized workloads  
Hardware for specialized workloads  
Short-lived compute instances  
Hardened VMs  
Dedicated physical servers  
VMware on Compute Engine

## STORAGE

Cloud Filestore  
Cloud Storage  
Persistent Disk  
Local SSD

Managed NFS server  
Multi-class multi-region object storage  
Block storage for VMs  
VM locally attached SSDs

## DATABASE

Cloud Bigtable  
Cloud Firestore  
Cloud Memorystore  
Cloud Spanner  
Cloud SQL  
Database Migration Service  
DB Insights

Petabyte-scale, low-latency, non-relational  
Serverless NoSQL document DB  
Managed Redis and Memcached  
Horizontally scalable relational DB  
Managed MySQL, PostgreSQL, SQL Server  
Migrate to Cloud SQL  
SQL Inspector

## DATA AND ANALYTICS

Query  
Query BI Engine  
Query ML  
Composer  
Data Fusion  
Dataflow  
Dataprep  
Dataproc  
Pub/Sub  
Data Catalog  
Data Studio  
Looker

Data warehouse/analytics  
In-memory analytics engine  
Query model training/serving  
Managed workflow orchestration service  
Graphically manage data pipelines  
Stream/batch data processing  
Visual data wrangling  
Managed Spark and Hadoop  
Global real-time messaging  
Metadata management service  
Collaborative data exploration/dashboarding  
Enterprise BI and Analytics

## HYBRID AND MULTI-CLOUD

Anthos  
Anthos Clusters  
Anthos Config Management  
Anthos Service Mesh  
Cloud Run for Anthos  
GCP Marketplace for Anthos  
Migrate for Anthos  
Operations  
Cloud Build  
Traffic Director  
Apigee API Management

Enterprise hybrid/multi-cloud platform  
Hybrid/on-prem Kubernetes Engine  
Policy and security automation  
Managed service mesh (Istio)  
Serverless development for Anthos  
Pre-configured containerized apps  
Migrate VMs to Kubernetes Engine  
Monitoring, logging, troubleshooting  
Continuous integration/delivery platform  
Service mesh traffic management  
API management, development, security

## API PLATFORM AND ECOSYSTEMS

API Analytics  
API Monetization  
Apigee API Platform  
API Gateway  
Apigee Hybrid  
Apigee Sense  
Cloud Endpoints  
Cloud Healthcare API  
Developer Portal  
GCP Marketplace  
AppSheet

API metrics  
Monetize APIs  
Develop, secure, monitor APIs  
Fully managed API Gateway  
Manage hybrid/multi-cloud API environments  
API protection from attacks  
Cloud API gateway  
Healthcare system GCP interoperability  
API management portal  
Partner & open source marketplace  
No-code App creation

## MAPS PLATFORM

Directions API  
Distance Matrix API  
Geocoding API  
Geolocation API  
Maps Embed API  
Maps JavaScript API  
Maps SDK for Android  
Maps SDK for iOS  
Maps Static API

Get directions between locations  
Multi-origin/destination travel times  
Convert address to/from coordinates  
Derive location without GPS  
Display iframe embedded maps  
Dynamic web maps  
Maps for Android apps  
Maps for iOS apps  
Display static map images

## AI/ML

AI Platform Data Labeling  
AI Platform Deep Learning VMs  
AI Platform Deep Learning Containers  
AI Platform Notebooks  
AI Platform Pipelines  
AI Platform Predictions  
AI Platform Training  
AI Platform  
AutoML, Natural Language  
AutoML Tables  
AutoML Translation  
AutoML Video Intelligence  
AutoML Vision  
Cloud Natural Language API  
Cloud Speech-To-Text API  
Cloud Talent Solutions API  
Cloud Text-To-Speech API  
Cloud TPU  
Cloud Translation API  
Cloud Video Intelligence API  
Cloud Vision API  
Contact Center AI  
Dialogflow  
Document AI  
Explainable AI  
Recommendations AI  
Vision Product Search

Data labeling by humans  
Preconfigured VMs for deep learning  
Preconfigured containers for deep learning  
Managed JupyterLab notebook instances  
Hosted ML workflows  
Autoscaled model serving  
Distributed AI training  
Managed platform for ML  
Custom text models  
Custom structured data models  
Custom domain-specific translation  
Custom video annotation models  
Custom image models  
Text parsing and analysis  
Convert audio to text  
Job search with ML  
Convert text to audio  
Hardware acceleration for ML  
Language detection and translation  
Scene-level video annotation  
Image recognition and classification  
AI in your contact center  
Create conversational interfaces  
Analyze, classify, search documents  
Understand ML model predictions  
Create custom recommendations  
Visual search for products

## NETWORKING

Carrier Peering  
Direct Peering  
Dedicated Interconnect  
Partner Interconnect  
Cloud Armor  
Cloud CDN  
Cloud DNS  
Cloud Load Balancing  
Cloud NAT  
Cloud Router  
Cloud VPN (HA)  
Network Service Tiers  
Network Telemetry  
Traffic Director  
Cloud Service Mesh  
Virtual Private Cloud  
VPC Service Controls  
Network Intelligence Center

Peer through a carrier  
Peer with GCP  
Dedicated private network connection  
Connect on-prem network to VPC  
DDoS protection and WAF  
Content delivery network  
Programmable DNS serving  
Multi-region load distribution/balancing  
Network address translation service  
VPC/on-prem network route exchange (BGP)  
VPN (Virtual private network connection)  
Price vs performance tiering  
Network telemetry service  
Service mesh traffic management  
Service-aware network management  
Software defined networking  
Security perimeters for API-based services  
Network monitoring and topology

## GAMING

Cloud Game Servers  
Orchestrate Agones clusters

## INTERNET OF THINGS (IOT)

Cloud IoT Core  
Manage devices, ingest data

## IDENTITY AND SECURITY

Access Transparency  
Assured Workloads  
Binary Authorization  
Certified Authority Service  
Cloud Asset Inventory  
Cloud Audit Logs  
Cloud DLP  
Cloud HSM  
Cloud EKM  
Cloud IAM  
Cloud Identity  
Cloud Identity-Aware Proxy  
Cloud KMS  
Cloud Resource Manager  
Security Command Center  
Cloud Security Scanner  
Confidential Computing  
Context-aware Access  
Event Threat Detection  
Managed Service for Microsoft Active Directory  
Secret Manager

Audit cloud provider access  
Workload compliance controls  
Kubernetes deploy-time security  
Managed private CA  
All assets, one place  
Audit trails for GCP  
Classify and redact sensitive data  
Hardware security module service  
External keys you control  
Resource access control  
Manage users, devices & apps  
Identity-based app access  
Hosted key management service  
Cloud project metadata management  
Security management & data risk platform  
App engine security scanner  
Encrypt data in-use  
End-user attribute-based access control  
Scans for suspicious activity  
Managed Microsoft Active Directory  
Store and manage secrets

## IDENTITY AND SECURITY (CONT.)

Security Key Enforcement  
Shielded VMs  
Titan Security Key  
VPC Service Controls

Two-step key verification  
Hardened VMs  
Two-factor authentication (2FA) device  
VPC data constraints

## MANAGEMENT TOOLS

Cloud APIs  
Cloud Billing API  
Cloud Billing  
Cloud Console  
Cloud Deployment Manager  
Cloud Mobile App  
Private Catalog  
Cloud Debugger  
Error Reporting  
Cloud Logging  
Cloud Monitoring  
Cloud Profiler  
Cloud Trace  
Transparent SLIs

APIs for cloud services  
Programmatically manage GCP billing  
Billing and cost management tools  
Web-based management console  
Templated infrastructure deployment  
iOS/Android GCP manager app  
Internal Solutions Catalog  
Live production debugging  
App error reporting  
Centralized logging  
Infrastructure and application monitoring  
CPU and heap profiling  
App performance insights  
Monitor GCP services

## DEVELOPER TOOLS

Cloud Build  
Cloud Code for IntelliJ  
Cloud Code for VS Code  
Cloud Code  
Cloud Tools for Eclipse  
Cloud Tools for Visual Studio  
Gradle App Engine Plugin  
Maven App Engine Plugin  
Cloud SDK  
Cloud Shell  
Artifact Registry  
Cloud Source Repositories  
Container Registry  
Container Analysis  
Cloud Scheduler  
Cloud Tasks  
Cloud Workflows

Continuous integration/delivery platform  
IntelliJ GCP tools  
VS Code GCP tools  
Cloud native IDE extensions  
Eclipse GCP tools  
Visual Studio GCP tools  
Gradle App Engine plugin  
Maven App Engine plugin  
CLI for GCP  
Browser-based terminal/CLI  
Universal package manager  
Hosted private git repos  
Private container registry/storage  
Automated security scanning  
Event-driven Cloud Run services  
Managed cron job service  
Asynchronous task execution  
HTTP services orchestration

## MIGRATION TO GCP

BigQuery Data Transfer Service  
Cloud Data Transfer  
Transfer Appliance  
Storage Transfer Service  
Migrate for Anthos  
Migrate for Compute Engine  
Migrate from Amazon Redshift  
Migrate from Teradata  
VM Migration  
Cloud Foundation Toolkit  
KF

Bulk import analytics data  
Data migration tools/CLI  
Rentable data transport box  
On-prem/on-premises data transfer  
Migrate VMs to GKE containers  
Compute Engine migration tools  
Migrate from Redshift to BigQuery  
Migrate from Teradata to BigQuery  
VM migration tools  
Infrastructure as Code templates  
Cloud Foundry to Kubernetes

## ADDITIONAL RESOURCES

Cloud Home Page  
Cloud Blog  
Cloud Platform Podcast  
Kubernetes Podcast  
Cloud Open Source  
GCP Medium Publication  
Apigee Blog  
Firebase Blog  
Workspace Developers Blog  
Workspace Github  
Workspace Twitter  
Cloud Certifications  
Cloud System Status  
Cloud Training  
Developers Blog  
Maps Platform Blog  
Open Source Blog  
Security Blog  
Kaggle Home Page  
Kubernetes Blog  
Regions and Network Map  
DORA - Software & Delivery Research

## MAPS PLATFORM (CONT.)

Maps SDK for Unity  
Shielded URLs  
Places API  
Places Library, Maps JS API  
Places SDK for Android  
Places SDK for iOS  
Roads API  
Street View Static API  
Street View Service  
Time Zone API

Unity SDK for games  
URL scheme for maps  
Rest-based Places features  
Places features for web  
Places features for Android  
Places feature for iOS  
Convert coordinates to roads  
Static street view images  
Street view for JavaScript  
Convert coordinates to timezone

## WORKSPACE

Admin SDK  
AMP for Email  
Apps Script  
Calendar API  
Classroom API  
Cloud Search  
Docs API  
Drive Activity API  
Drive API  
Drive Picker  
Email Markup  
Workspace Add-ons  
Workspace Marketplace  
mail API  
Hangouts Chat Bots  
People API  
Sheets API  
Slides API  
Task API  
Tasks API

Managed Workspace resources  
Dynamic interactive email  
Extend and automate everything  
Create and manage calendars  
Provision and manage classrooms  
Unified search for enterprise  
Create and edit documents  
Retrieve Drive activity  
Read and write files  
Drive file selection widget  
Interactive email using schema.org  
Extend Suite apps  
Storefront for integrated applications  
Enhance mail  
Conversational bots in chat  
Manage user's Contacts  
Read and write spreadsheets  
Create and edit presentations  
Search, read & update Tasks  
Manage your organization's eDiscovery

## MOBILE


Cloud Firestore  
Cloud Functions for Firebase  
Cloud Storage for Firebase  
Crashlytics  
A/B Testing  
App Distribution  
Authentication  
Cloud Messaging  
Dynamic Links  
Extensions  
Hosting  
In-App Messaging  
Performance Monitoring  
Predictions  
Realtime Database  
Remote Config  
Test Lab  
Analytics  
ML Kit

Document store and sync  
Event-driven serverless applications  
Object storage and serving  
Crash reporting and analytics  
Create A/B test experiments  
Trusted tester early access  
Drop-in authentication  
Send device notifications  
Link to app content  
Pre-packaged development solutions  
Web hosting with CDN/SSL  
Send in-app contextual messages  
App/web performance monitoring  
Predict user targeting  
Real-time data synchronization  
Remotely configure installed apps  
Mobile testing device farm  
Mobile app analytics  
ML APIs for mobile

# Dynamic Datacenter Infrastructure

## Legacy Infrastructure

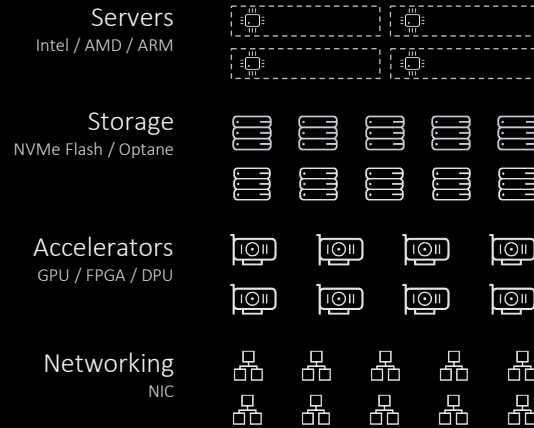
Static Servers



Statically Configured

## Modern Infrastructure

Disaggregated Resource Pools




Dynamically Configurable Software Defined Hardware

Dynamically Configured Servers




# Use Case: GPU Workload Sharing

Daytime Workload




7am to 4pm

GPU Pool



Radically Improve GPU Utilization

Nighttime Workload




4pm to 7am



# Use Case: GPU On-Demand/HPC Cloud Provisioning


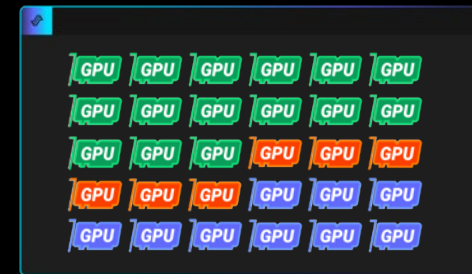
## Statically Configured GPUs



- Limited Scalability
- Inflexible
- Poor Resource Utilization

## Dynamic Configurable GPUs

### Composable GPU On-Demand



- High GPU Density
- Rapid Deployment and Scaling
- Increased Resource Utilization



- Recent Higher Ed Installation
  - 24 Dell C6525 servers (2U/ 4 node)
  - 56 A100 80GB GPUs
    - Expandable for additional GPUs
  - Liquid IO Accelerators w Intel Optane PS
- Recent DoD installations
  - Three HPC clusters
  - More than 2000 compute nodes
  - 892 Nvidia A100 GPUs

Where do we go from here?

The image features a large, stylized logo for CXL (Compute Express Link) in shades of teal and light green. A white diagonal slash is drawn through the logo. To the right of the slash, the text 'Compute Express Link™' is displayed in a smaller, grey font, with the 'X' in 'Express' matching the light green color of the main logo.

# Compute Express Link

Rapid expansion of cloud computing

Growth of AI and analytics

Industry demands faster data processing, higher performance, lower latency

# Compute Express Link

New class of CPU to device  
interconnect for heterogenous systems

- PCIe Gen5 physical & electrical paths
  - CPU/Memory/Device sharing
- Very tight latency tolerances for disaggregated memory

# Compute Express Link

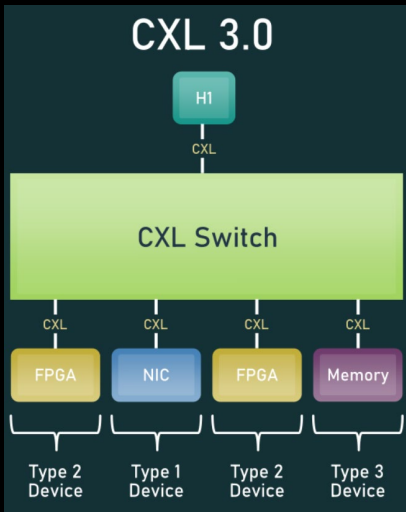
Features	CXL 1.0 / 1.1	CXL 2.0	CXL 3.0
Release date	2019	2020	1H 2022
Max link rate	32GTs	32GTs	64GTs
Flit 68 byte (up to 32 GTs)	✓	✓	✓
Flit 256 byte (up to 64 GTs)			✓
Type 1, Type 2 and Type 3 Devices	✓	✓	✓
Memory Pooling w/ MLDs		✓	✓
Global Persistent Flush		✓	✓
CXL IDE		✓	✓
Switching (Single-level)		✓	✓
Switching (Multi-level)			✓
Direct memory access for peer-to-peer			✓
Enhanced coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			✓
Multiple Type 1/Type 2 devices per root port			✓
Fabric capabilities (256 byte flit)			✓

CXL Consortium 2022

# CXL 3.0 – *The building block for disaggregation*

“Modern datacenters require heterogenous and composable architectures to support compute intensive workloads for applications such as Artificial Intelligence and Machine Learning – and we continue to evolve CXL technology to meet industry requirements...”

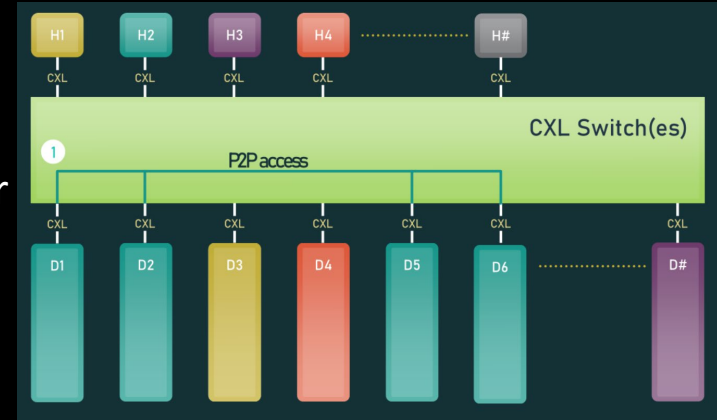
*Siamak Tavallaei, CXL Consortium President Aug 2022*



Multi-device to host

CXL Consortium 2022

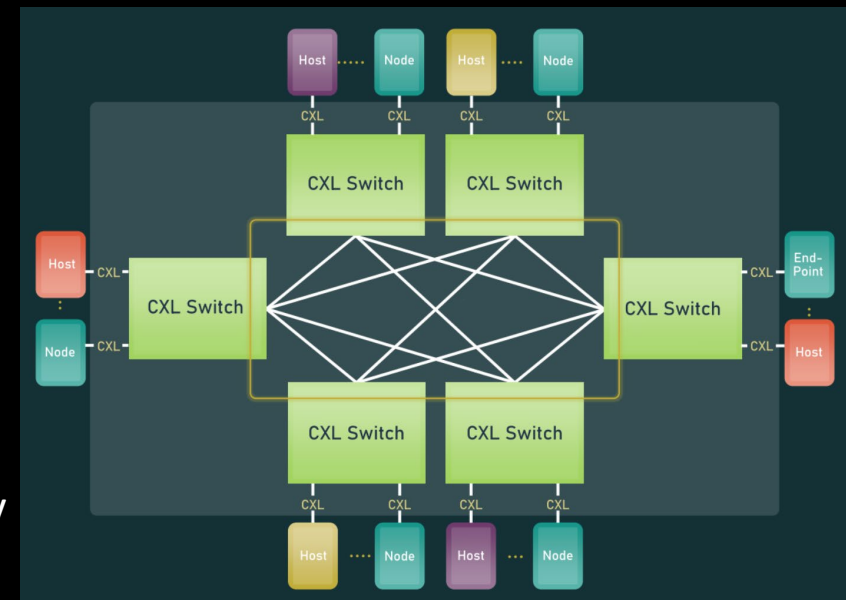
Device Peer 2 Peer memory sharing



CXL Consortium 2022

Node agnostic fabric connections

- Hosts
- Devices w Cache
- Devices w Cache + Memory
- Devices w Memory



CXL Consortium 2022

Thank You

