CYGNA LABS DDI for Zero Trust Network Architectures Owner's Manual

CYGNA LABS

White Paper

DDI for Zero Trust Network Architectures

By Timothy Rooney

Contents [hide]

Introduction

The burgeoning adoption of cloud applications and networking, increasing quantities and types of network-connected devices, and sprawling network domains not only into cloud services but to Internet of Things (IoT) deployments and to remote and mobile workers, have led to skyrocketing complexity in enterprise network topologies over the last decade. Network evolution has not only challenged network managers to keep up with this blistering pace of change, it has exposed a vastly larger attack surface from a network security perspective. With workers requiring access to cloud applications from wherever they happen to be using a company or personal device, the former perimeter-based security model of partitioning security for internal vs. external networks, vigilantly linked through one or more demilitarized zone (DMZ) firewalls has been summarily supplanted by ubiquitous Internet access from anywhere.

Localized Internet access, referred to as Internet breakout in SD-WAN parlance, enables a device to connect to its geographically closest cloud application hosting site instead of traversing the enterprise network to a DMZ perimeter to connect. Internet break- out shortens the path to the cloud, reducing latency and optimizing application performance. But the dissolution of a well-protect- ed portcullis through which Internet traffic could be surveyed and blocked leads to a larger quantity of potentially less defended entrée points to the enterprise network from the Internet.

The concept of zero trust embraces this new reality and disposes of any implicit trust of a user or device based on internal vs. external source IP address, known MAC address, prior login history, or legitimate access grants to other applications. Every attempt to access an enterprise resource must be authenticated and authorized. The logical entities through which the zero trust architecture performs this authentication and authorization functions are the policy decision point (PDP) and the policy enforce- ment point (PEP).

This white paper provides an introduction to zero trust, and focuses on the role of network-foundational DHCP-DNS-IPAM (DDI) services in a zero trust deployment. While DDI services are not implicitly in-band like a PDP/PEP gateway that can permit or deny traversal network traffic, they do serve as indispensable helper services, providing automated IP address assignment and user and device network navigation. DDI

services are instrumental in detecting and protecting networks from a security perspective in general, and we'll explore how they apply to zero trust in particular.

The U.S. Government Zero Trust Mandate

The United States recently declared a strategy for all Federal agencies to implement a zero trust architecture. This strategy was issued by the Office of Management and Budget (OMB) in January, 2022 in support of a May, 2021 Executive Order (14028) to improve the nation's cybersecurity. The strategy sets forth five key goals for agencies to implement:

1. **Identity:** Staff within each agency must use managed identities to access the applications they use in their work. Phishing- resistant multi-factor authentication (MFA) protections should be implemented.
2. **Devices**: The Federal Government shall maintain a complete inventory of every device it operates and authorizes for its use, and can prevent, detect, and respond to incidents on those devices.
3. **Networks**: Agencies shall encrypt all DNS requests and HTTP traffic within their environment and begin executing a plan to shrink their perimeters into isolated environments.
4. **Applications and Workloads**: All applications shall be treated as Internet-connected, and agencies shall routinely subject their applications to rigorous empirical testing, and welcome external vulnerability reports.
5. **Data**: Agencies shall deploy protections that make use of thorough data categorization, take advantage of cloud security services to monitor access to their sensitive data, and implement enterprise-wide logging and information sharing.

In order of DDI implication these five goals could be ordered Networks, Devices, Data, Applications and Workloads, and Identity. As one of the D's in DDI, encrypting DNS transactions called for in the Networks goal falls squarely within the realm of DDI. The shrinking of perimeters also affects the DDI IPAM function from an IP addressing/subnetting perspective as resources with a network micro-perimeter may likely reside on a common subnet if not VLAN. We'll delve more deeply into the DDI imper- atives for each of these goals later.

Device inventory within the Devices goal should include current IP address assignment as well as IP address assignment history for mobile devices or renumbered networks. Monitoring access to sensitive data called for within the Data goal should include monitoring DNS queries as one component of this goal. As the first step in nearly every network connection attempt, whether legitimate or by malware, DNS directs the querier to an IP address to which to connect, so tracking DNS transactions can provide valuable cyber threat information.

Tracking of vulnerability reports and empirical testing should include DDI components, including appliance kernel or operating systems vulnerabilities, as well as those for DHCP, DNS and IPAM applications. Use of managed identities minimally affects DDI, though machine-initiated authentication processes may utilize device name and/or IP address assignment information.

These goals extend beyond the US government to other organizations seeking to implement zero trust, as zero trust is also a core component of emerging network architectures such as Secure Access Service Edge (SASE). So like the National Institute of Standards and Technology (NIST) Cybersecurity Framework, this zero trust framework applies to most enterprises globally.

Zero Trust Overview

While zero trust is among the leading emergent security approaches, its origins can be traced back over twenty years ago with the publication of RFC 2753 1 , “A Framework for Policy-based Admission Control,” though this RFC proposed such a framework for quality-of-service allocations. Nevertheless, the concept of a policy enforcement point (PEP) and policy decision point (PDP) used in zero trust architectures was described in this RFC published in January 2000.

The concept of zero trust networks, originally so named by Forrester Research a decade ago, is rising in prominence as a fun- damental network security approach. As the name implies, zero trust networks begin with the assumption that no user or de- vice is implicitly trusted. Contrast this with the former castle-and-moat philosophy where users and devices within a network were implicitly trusted and defenses focused on detecting and repelling attacks originating externally to the network. Users’ demands for access to

networks from a growing diversity of computing devices with high mobility has raised the network's vulnerability to attacks originating within the network, maliciously or otherwise.

In advance of the US mandating implementation of a zero trust architecture (ZTA) across Federal agencies, NIST had released special publication 800-207 2 ("NIST SP 800-207"), Zero Trust Architecture. This document defines seven tenets of zero trust, which are positioned as considerations more so than hard requirements.

1. All data sources and computing services are considered resources. This includes user devices, both enterprise-owned and us- er-owned, cloud systems, IoT devices, etc.
2. All communication is secured regardless of network location. Trust should not automatically be granted based on the use of an enterprise (e.g., private) IP address.
3. Access to individual enterprise resources is granted on a per-session basis. Trust in the requester is evaluated, and if granted, least privilege access is granted. Access approval to one resource though does not automatically grant access to a different resource.
4. Access to resources is determined by dynamic policy. The policy is based on the observable state of client identity, application/ service, and the requesting asset.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorization are dynamic and strictly enforced before access is granted.
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications and uses it to improve its security posture.

The fundamental goal of zero trust boils down to preventing unauthorized access to data and services while applying granular access control enforcement as possible.

Zero Trust Architecture

Zero trust architectures seek to achieve these goals of preventing unauthorized access while applying least-privilege access on a granular level with policy decision point (PDP) and policy enforcement point (PEP) logical entities. Following a modern "software-defined" architecture with the PDP residing in the control plane and the PEPs in the data

plane enables the centralized declaration of access policies with distributed enforcement.

Figure 1 illustrates zero trust logical components. The PDP consists of a policy engine and a policy administrator. The policy engine leverages enterprise-defined access policies along with current network, compute, personnel, and application information to derive real time access decisions for new and existing connections. This information includes monitoring (continuous diagnostics and mitigation, CDM) system events, industry compliance requirements, threat intelligence, activity and audit logs, organizational PKI (e.g., certificate) information, identity management updates, and SIEM events. The policy engine processes these various inputs to define or modify existing access dispositions. The policy administrator component of the PDP disseminates these access decisions to relevant PEPs throughout the network.

The PEP is responsible for enabling, monitoring, and terminating connections between one enterprise resource and another. As depicted in Figure 1, the PEP sits logically between a subject attempting to access, via a system, an enterprise resource. The PEP is essentially a gateway that applies access decisions received from the PDP to allow or deny the intended connection. A connection may be allowed or denied upon initial request, or an existing connection may be terminated if updated information from the PDP indicates denial of access.

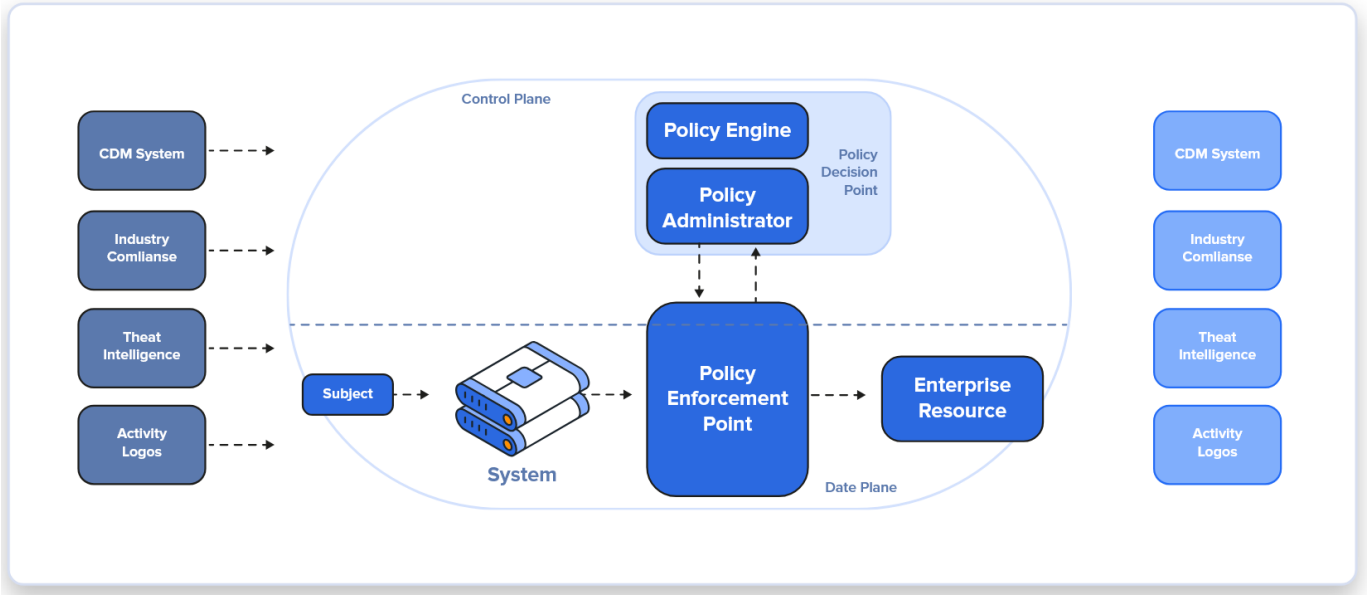

Figure 1: Zero Trust Logical Components

Source: Zero Trust Architecture, NIST, SP-800-207, p. 9.

Zero Trust Implementation Approaches

SP-800-207 defines three basic architectural variants when implementing zero trust, though they would likely be combined within an enterprise's overall approach. The first approach emphasizes identity governance, focused on the subject (of Figure 1) attempting to access a given resource. Each subject's identity must be authenticated, then least-privilege authorization is applied to the use of the corresponding resource.

The second variant applies network segmentation to deploy individual or sets of resources on protected network segments. Referring to Figure 1, the PEP serves as the gateway to the protected or trusted segment on which resides an enclave of resources for example. A more stringent application could place the logical PEP on individual device hosts, applying access controls at the host level.

The third approach utilizes network infrastructure and software-defined perimeters. One incarnation of this approach applies an overlay at the application [or a lower OSI] layer with an agent/gateway implementation where the client system uses an installed agent that establishes a secure connection to the associated resource. In this scenario, the PEP is logically split be- tween the agent, deployed on or near the requesting system, and the resource or a nearby proxy to the resource.

Your zero trust implementation may likely utilize all three variants for different resource types or for different topologies e.g., cloud-based, partner access, etc. In general, implementing zero trust requires the identification of your assets (virtual, physical and possibly user devices), subjects (including user privileges), and business processes (workflows and data flows).

Zero Trust Maturity Model

The Zero Trust Maturity Model was created by the Cybersecurity and Infrastructure Security Agency (CISA) of the U.S. government and defines one of many paths to designing and implementing zero trust within an organization. As such, it serves as a credible guide for planning your zero trust deployment, while acknowledging, "is an incremental process that may take years to implement." The maturity model defines a

continuum of zero trust maturity along five pillars, which map to the five zero trust goals outlined in Executive Order 14028, namely Identity, Devices, Networks, Applications and Workloads, and Data.

This continuum cuts across each of the five pillars, and is segmented into four grades of maturity, starting from Traditional, moving up through Initial and Advanced levels, ultimately to the Optimal maturity level. These levels are illustrated in Figure 2 and are defined as follows:

- **Traditional** – Largely manual and static security policies and solutions that address one pillar at a time with siloed enforcement of policies.
- **Initial** – Some automation for element and security policy configuration and enforcement with initial cross-pillar coordination and adaptation of policies post-provisioning.
- **Advanced** – Centralized visibility and control with automation for network and security policy configuration and enforcement with automated responses with pre-defined mitigations with adaptation to policies with rich cross-pillar coordination, building towards enterprise-wide awareness.
- **Optimal** – Fully automated, just-in-time policy configuration for assets and resources that dynamically self-manage policies based on observed triggers with cross-pillar interoperability with continuous monitoring, centralized visibility, and comprehensive situational awareness.

The maturity model defines several criteria for each maturity level for each pillar that serves to guide an organization to understand their current stance with respect to zero trust maturity, and potential next steps to advance to higher maturity levels across the five pillars.

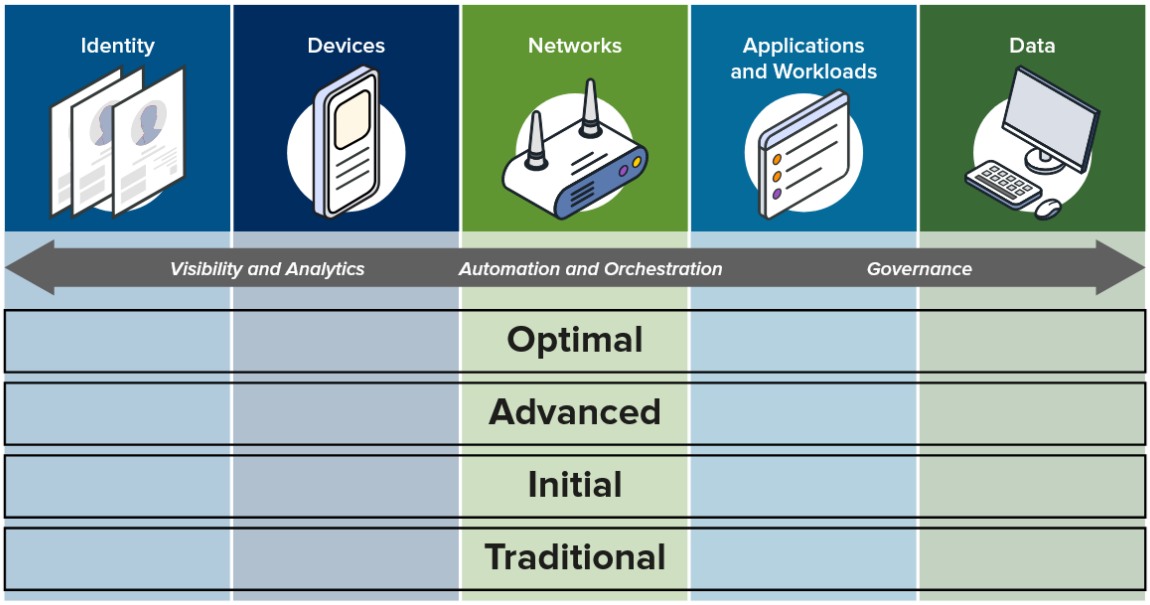

Figure 2: Zero Trust Maturity Model framework
Source: Derived from Zero Trust Maturity Model, Version 2.0, CISA, April 2023, p.10.

Each pillar also highlights cross-cutting capabilities, which support interoperability across the five pillars, defined as follows:

- **Visibility and Analytics** – Visibility refers to observable events and transactions across the network and analytics relates to processing observed transactions or events over time or across multiple sources across the network to help inform policy decisions, proactive measures, and facilitate responses.
- **Automation and Orchestration** – Automated tools and workflows to support security response functions across products and services while maintaining oversight and security.
- **Governance** – Definition and enforcement of cybersecurity policies, procedures, and processes within and across pillars to manage and mitigate security risks in support of zero trust principles.

Zero Trust DDI Imperatives

The Zero Trust Maturity Model specifies proficiency characteristics for each maturity level and cross-cutting capabilities within each pillar. It provides a means to identify aspects of desired maturity to define incremental zero trust implementation steps. Given DDI is

foundational to network architectures and initiatives, let's examine the impacts of each of these pillars from a DDI perspective. We define DDI as comprised of the following components:

- **DNS (Domain Name System)** is a core network infrastructure service that serves as the first step in making a network connection by translating host names like web addresses into IP addresses among other things. As such it also serves as a key component for network security and cyberthreat investigations.
- **DHCP (Dynamic Host Configuration Protocol)** automates the assignment of network configurations, including relevant IP addresses for IPv4 and IPv6, to network devices in an unattended fashion.
- **IPAM (IP Address Management)** efficiently manages the IP address foundation within networks, cloud systems, Internet of Things (IoT) and related deployments, and supports DNS and DHCP technologies.

In summary, DDI solutions integrate DNS, DHCP and IPAM, thus combining network and cloud-based IP address configurations while facilitating the secure operation of core network services. DDI technologies are essential for effective network operation. Contrasting with the five zero trust pillars, derivative DDI imperatives include:

- **Networks**
  - **Encryption** – Zero Trust requires encryption of network traffic; encryption of DHCP traffic is undefined given the net- work-detached state of a device seeking configuration via the DHCP protocol. However, DNS standards exist to encrypt DNS queries and responses via transport layer security (TLS) and over HTTPS, referred to as DNS over TLS (DoT) and DNS over HTTPS (DoH) respectively.
  - **Network segmentation** – creation of micro-perimeters, which shrinks the trust perimeter from the "entire enterprise network" to a constrained segment of resources.
  - **Network inventory** – maintenance of a plan of record or source of truth for IP networks, subnets, address pools, and individual IP address assignments is imperative as a means to define the scope of the network. Visibility to new devices attempting to access the network via DHCP enables evaluation and assessment of the device.

- **Devices**
  - **Inventory** – Retaining a device inventory enables an organization to enumerate the full scope of its network and security focus. While DDI systems do not serve as full-on device inventory systems, they do maintain a centralized source of truth for device IP address and domain details.
  - **Discovery** – Dynamic tracking of device DDI inventory via DHCP, SNMP or ping discovery, IoT border router neighbor table discovery or cloud API fetches keeps the device inventory refreshed and up to date in reflecting the actual state of device presences in the networks.
- **Data**
  - **Monitoring** – Tracking network activity provides visibility and detection of potential nefarious activities. From a DDI perspective, monitoring DNS, as the first step in an IP connection, supplements a defense in depth network security stance.
  - **Logging** – aggregation of DDI data for audits, spot checks, reporting, or forensics is indispensable.
- **Applications and workloads**
  - **Paths** – tracking of network layer endpoint parameters provides input to paths from subjects to applications and workloads for identification of and enforcement of network micro perimeters.
  - **Visibility** – monitoring not only endpoints but traffic flows within serves as input to illicit activity such as malware lateral movement or data exfiltration.
- **Identity**
  - **DDI system identity** – DDI manages network layer configurations, changes, and events, offering minimal input to subject identity details other than IP address and DNS information. Access to DDI systems however and scoped administrator least-privilege access is critical to constraining domains of control beyond job function requirements.

Cygna DDI Support for Zero Trust

Considering these DDI imperatives, the zero trust maturity model, and three basic implementation approaches discussed earlier, your DDI solution plays a key role in your zero trust deployment initiatives. As a leading DDI vendor, Cygna Labs offers a variety of

capabilities to support your overall zero trust implementation. Zero trust DDI imperatives are addressed with our products in the following ways:

- **Networks**
  - **Encryption** – Zero Trust requires encryption of network traffic. Cygna DDI solutions support transport layer security (TLS) communications among system components and HTTPS administrator and REST access. Our DNS services support DoH and DoT as well.
  - **Network segmentation** – Cygna DDI solutions excel at defining multi-layered IP address block hierarchies that enable IT administrators to define network allocations to regions, sites, and subnets, VLANs, then down to micro-perimeters, e.g., a segment with a PEP protecting a segment of resources.
  - **Network inventory** – Cygna DDI solutions serve as the plan of record or source of truth for your IP networks, subnets, address pools, and individual IP address assignments. Visibility to new devices attempting to access the network via DHCP and network discovery features enable reconciliation of the DDI inventory with the network.
- **Devices**
  - **Inventory** –While our DDI systems do not serve as full-on device inventory systems, they do maintain a centralized source of truth for device IP address and domain details.
  - **Discovery** – Cygna Labs' DDI solutions provide dynamic tracking of device DDI inventory via DHCP, SNMP or ping discovery, IoT border router neighbor table discovery and cloud API to maintain accuracy of the device inventory to reflect the actual state of device IP/DNS information in the networks.
- **Data**
  - **Monitoring** – Cygna DDI Guard collects and archives DHCP and DNS packets as well as DNS tunneling detection events and domain generation algorithm (DGA) DNS queries.
  - **Logging** – Logging to generic SIEM systems is supported, though DDI-centric DDI Guard provides built-in DDI reports and statistics for analysis and cyberthreat investigations. DDI Guard also supports filtering of unremarkable DHCP and DNS in- formation (e.g., queries to well-known domains) in transmitting data to SIEM systems to reduce SIEM ingest and processing costs.

- **Applications and workloads**
  - **Paths** – DDI services are used to obtain IP addresses and resolve IP addresses and are not “in band” within the data path of connection endpoints. However, the endpoints may initially be defined through knowledge of a given devices’ IP address and DNS queries and responses from/to that device. This information can be helpful during CTI to determine how or why a given connection, detected through an in-band device like a firewall, occurred.
  - **Visibility** – Cygna DDI Guard offers centralized visibility to DNS and DHCP packet-level transactions across your network, providing dashboards, reports, and alerts for anomalous conditions. An archive system is also available to store more data over time to address retention and compliance requirements. IPAM Auditor likewise provides insight into DNS queries and responses and DHCP with rich graphical dashboards and reports, and packet level drill-down for forensics. Both solutions provide these features plus drill-down to individual DNS queries and responses to enrich CTI.
- **Identity**
  - **DDI system identity** – Cygna DDI systems provide external authentication, 2FA, SAML, and OAuth2 support for integration with centralized identity management solutions.

In terms of the cross-cutting capabilities, Cygna Labs also supports corresponding DDI capabilities as follows:

- **Visibility and Analytics** – Cygna DDI Guard provides full visibility to DDI transactions observed by an organizations DHCP and DNS servers. This information is available for reporting and analytics to identify potential threats.
- **Automation and Orchestration** – REST APIs and callouts/user exits provide a means to invoke external systems due to a DDI detected event. Our Cygna Automation Appliance (CAA) enables centralized orchestration of DDI workflows with a simple graphical drag-and-drop workflow user interface. It enables the creation of REST endpoints, which invoke corresponding workflows to automate DDI related activities.
- **Governance** – Cygna DDI solutions feature centralized control-plane management architectures to promote pan-network visibility and application of DDI governance policies.

Summary

Deploying zero trust requires thorough network and application flow analysis and diligent application of corresponding zero trust measures. Several US Government standards and documents are available to provide guidance, including the Zero Trust Maturity Model, which enables organizations to gauge their current level of zero trust maturity and to identify next steps toward advancing toward optimal maturity. As the foundation of your network, DDI services perform a pivotal role in granting network access and directing application connectivity. Cygna Labs offers your choice of diverse DDI products and services to help you plan for and implement zero trust controls at the critical DDI layer.

About Cygna Labs

Cygna Labs is a software developer and one of the top three global DDI vendors. Many Fortune 100 customers rely on Cygna Labs' DDI products and services, in addition to its industry-leading security and compliance solutions to detect and proactively mitigate data security threats, affordably pass compliance audits, and increase the productivity of their IT departments. For more information, visit **cygnalabs.com**.

Cygna Labs Corp

sales@cygnalabs.com
Toll Free: 844.442.9462
Intl: +1 305-501-2430
www.cygnalabs.com

https://qrco.de/bfvBwn

Documents / Resources

	CYGNA LABS DDI for Zero Trust Network Architectures [pdf] Owner's Manual DDI-2025, DDI for Zero Trust Network Architectures, Zero Trust Network Architectures, Trust Network Architectures, Network Architectures, Architectures

References

- User Manual

CYGNA LABS

Architectures, CYGNA LABS, DDI for Zero Trust Network Architectures, DDI-2025, Network Architectures, Trust Network Architectures, Zero Trust Network Architectures

Leave a comment

Your email address will not be published. Required fields are marked *

Comment *

Name

Email

Website

Save my name, email, and website in this browser for the next time I comment.

Post Comment

Search:

e.g. whirlpool wrf535swhz

Search