



SOLUTION BRIEF

INTER-VSAN ROUTING WITH THE CISCO MDS 9000 FAMILY OF SWITCHES AND CISCO SAN-OS 2.1

This paper describes the enhancements to Inter-VSAN Routing (IVR) in the Cisco® SAN-OS 2.1.

SUMMARY

In today's complex SANs, scaling while maintaining a high level of availability is the most difficult problem for any SAN architect. The combination of the award-winning hardware and industry-leading software features offered by Cisco Systems® has enabled the company to bring a level of scalability and availability that is unparalleled in the industry.

Cisco Virtual SAN (VSAN) technology has changed the way SANs are being deployed. It has given customers design flexibility previously available only in more developed network technologies such as Ethernet. The VSAN concept has made such an impact on the industry that it has been adopted by ANSI T11 as the standard by which all virtual fabric implementations should be based.

IVR was a natural evolution of that same VSAN technology. By routing between VSANs, devices can maintain the level of separation in terms of fabric services and fabric wide events required for the highest level of availability yet take advantage of data sharing across thousands of devices. With the enhancements to IVR available in the Cisco SAN-OS 2.1, Cisco has extended its technology leadership in SAN routing, scalability, and availability. Two important enhancements to IVR in the Cisco SAN-OS 2.1—Fibre Channel ID (FCID) Network Address Translation (NAT) and enhanced scalability—make IVR the most flexible Fibre Channel routing solution available today.

IVR NAT

In Fibre Channel, device addressing is handled in two ways. The first, and more commonly visible to the end user, is the worldwide name (WWN) of a device. This 64-bit address uniquely identifies each device globally to ensure no duplicate WWNs in the Fibre Channel network. This is commonly used to make basic user-level management changes such as zoning for device access. The second scheme, much less visible to the end user, is FCID. The 24-bit FCID is designed to be a dynamic address assigned by the fabric when a device logs in to reduce complexity of addressing for internal use by the fabric. It comprises three components:

- **Domain**—The domain is a unique number assigned to each switch in a logical fabric. A domain ID assigned to a switch can range from 1 to 239. This number comprises the first 8 bits of the FCID.
- **Area**—The 8-bit area field is assigned by the switch as well. It can range from 0 to 256. In some third-party switches this number is assigned by using the physical port number (that is, port 3 out of 16 ports), limiting availability on some operating systems. The Cisco MDS assigns these sequentially regardless of the physical port number.
- **Port**—The port field is also 8 bits ranging from 0 to 256. This field is unique in that it also is used to assign the arbitrated loop physical address (ALPA) for devices that use loop. In the context of a device that is not using arbitrated loop, it is common to see the field set to 0, although this is not required.


Using these three fields, each device is assigned an FCID when it logs into the fabric. Comprising the 24 bits of the domain, area, and port, the FCID is used as a simplified addressing scheme, replacing the WWN internally to the fabric for everything from name server queries to routing. Although WWNs are unique globally, FCIDs are required to be unique only within the logical fabric in which they are assigned.

When SANs were first deployed, the domain field seemed sufficiently wide. With support for up to 239 switches in a single fabric, the domain field seemed enormous at a time when customers were deploying fabrics of 1 or 2 switches. Today, the scale of SAN design has changed dramatically.

For SANs potentially containing tens of thousands of ports, the 239 domain limitation does not seem so large. As more SAN consolidation occurs, reducing the number of physical networks while increasing the number of switches in these networks, the problem of domain ID assignment arises.

In many environments domain IDs already overlap in different physical infrastructures. This may be due to the large size of the SAN, miscalculation of how fast SANs would expand in the environment, or no plan to consolidate in the future. In many cases SAN islands were deployed (a common practice in the early SAN days). This meant that domain ID considerations were local to only to each island, not requiring planning for consolidation of domains. In any case, the problem is the same. When SAN consolidation starts and switches are interconnected, it is possible for two or more devices that must communicate to have the same domain ID (refer to Figure 1).

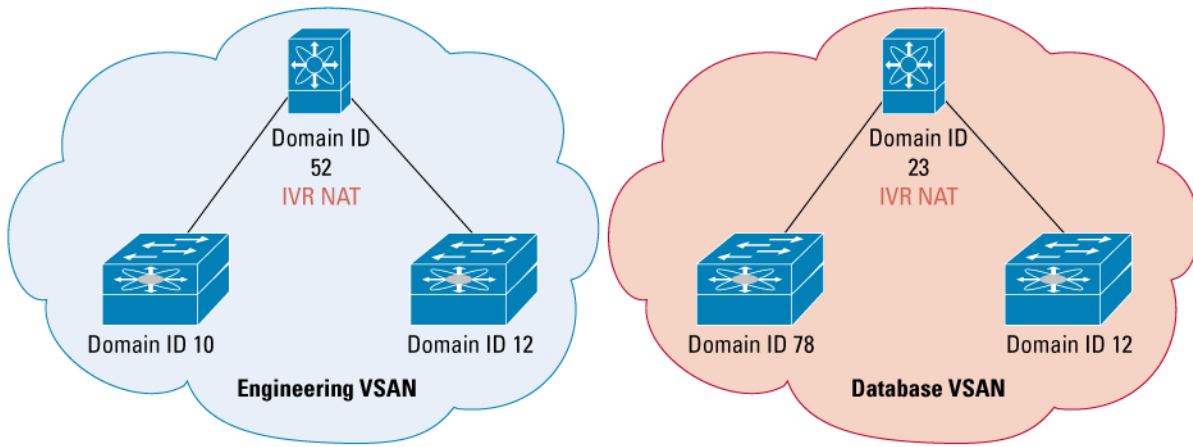
Figure 1. Independent SAN Islands can Have Overlapping Domain IDs.



In smaller networks, with fewer than 239 switches, in some cases the domain IDs may need to be changed when switches (with same domain ID) are connected to merge the network. When a domain ID changes, new FCIDs must be assigned to each device (because the domain ID comprises the first 8 bits of the FCID). This is a disruptive process, interrupting communication between the host and the disk for a short period of time. For some operating systems, such as AIX and HP-UX, a change of the FCID can cause an even larger impact (they may lose binding to their disk, requiring reestablishment at the OS level). In larger networks (with more than 239 switches), however, overlapping domain IDs across different logical fabrics (VSANs) is common because address is unique per logical fabric and may be exhausted in any given fabric.

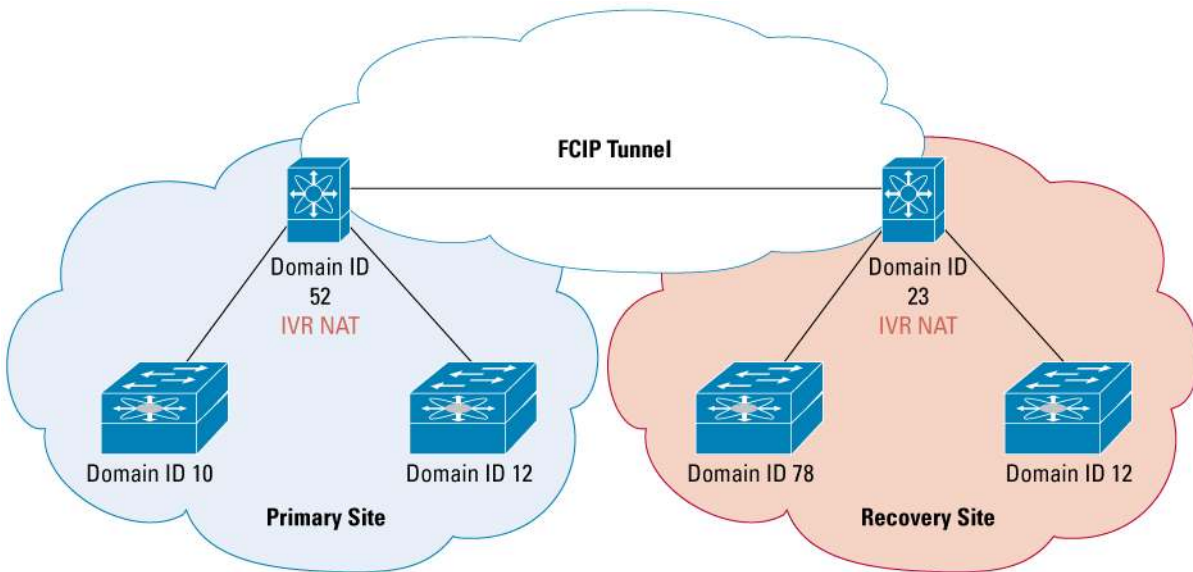
With IVR NAT, however, overlapping domain IDs across different VSANs is not a concern. As shown in Figure 2, when VSAN technology and IVR NAT are used, two logical fabrics with the same domain ID can share devices.

Figure 2. Inter VSAN Routing Handles the Translation Between Different VSANs with Overlapping Domains



This technology allows more scalability within a data center, and it also can be extended across long distances using either metropolitan-area network (MAN) technology or WAN technology with Fibre Channel over IP (FCIP). This allows the SAN to scale far beyond the local reach without having to maintain a globally unique addressing scheme (refer to Figure 3).

Figure 3. With IVR, NAT functionality can be used within the data center or extended across distances with FCIP or other SAN extension technology

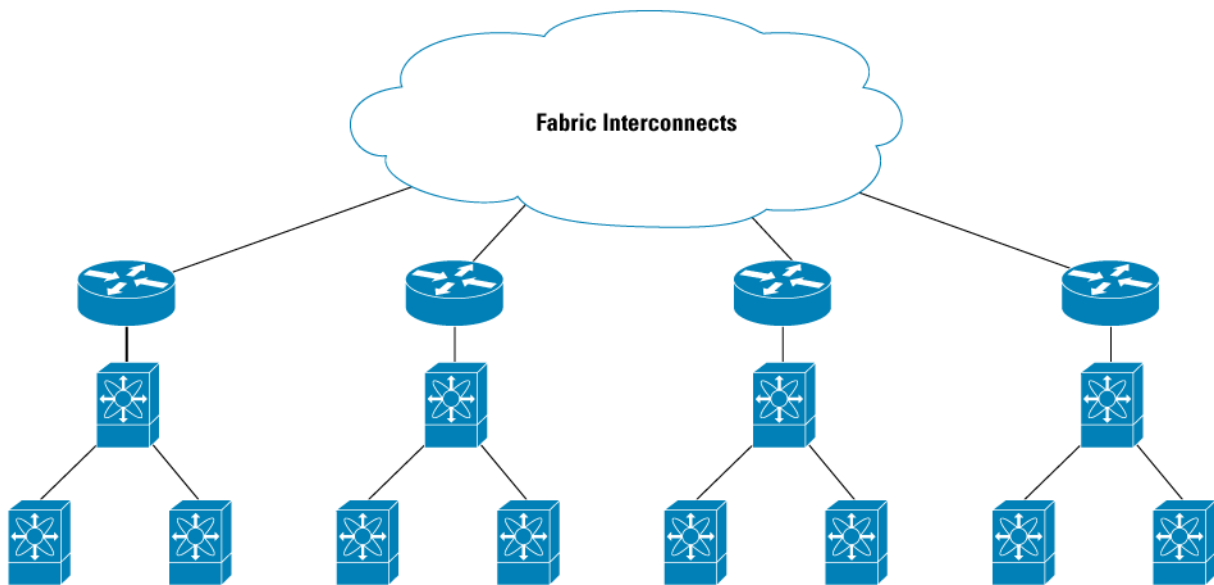


IVR with NAT allows the SAN consolidation required for today's business without requiring reconfiguration of the existing infrastructure. The requirement to change domain IDs—and therefore disrupt host-to-disk communication—is no longer necessary. The flexibility IVR NAT is available at no performance cost to the network. NAT functionality is handled in hardware by dedicated ASICs with no performance impact to mission critical application. As it did in IP networking, IVR NAT allows the flexibility to scale beyond the limitations of the available address space of Fibre Channel technology today.

IVR SCALABILITY

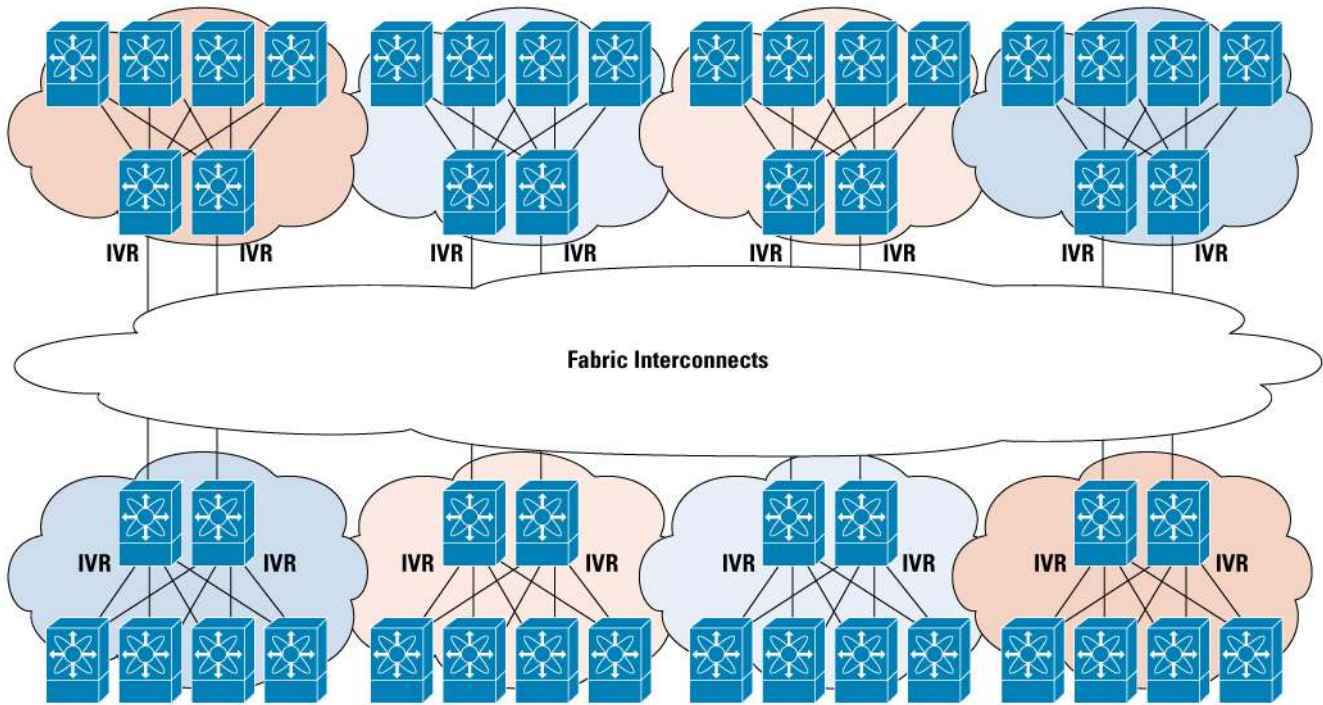
With the limit of switches in a physical infrastructure overcome by the use of NAT, the number of routes supported in the infrastructure needs to scale with it. In the infancy of routing technology, this meant adding dozens or even hundreds of external appliances to the network, adding complexity of operation and management and reducing overall availability in the environment. With today's advanced switching technology, hardware-based routing functions can be embedded directly into switches and directors, allowing for robust routing scalability without the need for external appliances (refer to Figure 4).

Figure 4. Legacy Networks use External Appliances for Routing Introducing Additional Latency, Management Requirements, and Reducing HA



IVR with the Cisco SAN-OS 2.1 enhances the scalability of routing within a single Cisco Fibre Channel switch or director to levels achievable only by a dozen or more external appliances. This integrated approach to routing scalability allows unparalleled support for today's largest and most complex SANs (refer to Figure 5).

Figure 5. With IVR in Cisco SAN-OS 2.1 SANs with Thousands of Devices can Have a High Level of Isolation with VSANs While Still Allowing Communication with Any Other VSAN



Fabric routing within the Cisco SAN-OS 2.1 can scale to thousands of routes per switch without the need for external routers. With support for 10000 devices across 128 VSANs IVR allows the most demanding connectivity requirements to be met without sacrificing availability or ease of management. This level of scalability helps enable greater consolidation of server, storage, and network resources, improving usage and decreasing costs throughout the data center.

CONCLUSION

The Cisco Fibre Channel fabric routing capability leads the industry in functions and scalability. By integrating routing within the switch rather than with an external appliance, the Cisco solution provides greater value and availability while reducing cost and complexity. With the ability to route between both Cisco switches and third-party switches, IVR also is the most interoperable solution in the industry. Building on this leadership position, the Cisco SAN-OS 2.1 helps enable even greater scalability and functions—with support for additional routes and NAT—allowing customers to achieve additional efficiencies and cost savings within their storage environments.

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R) 205252.m_ETMG_DB_4.05

